

THE HURWITZ SPACE OF COVERS OF AN ELLIPTIC CURVE E AND THE SEVERI VARIETY OF CURVES IN $E \times \mathbb{P}^1$

GABRIEL BUJOKAS

ABSTRACT. We describe the hyperplane sections of the Severi variety of curves in $E \times \mathbb{P}^1$ in a similar fashion to Caporaso–Harris’ seminal work [CH98]. From this description we *almost* get a recursive formula for the Severi degrees—we get the terms, but not the coefficients.

As an application, we determine the components of the Hurwitz space of simply branched covers of a genus one curve. In return, we use this characterization to describe the components of the Severi variety of curves in $E \times \mathbb{P}^1$, in a restricted range of degrees.

CONTENTS

1. Introduction	2
1.1. Hurwitz spaces	2
1.2. Severi varieties	5
1.3. Guide to readers	11
1.4. Acknowledgments	12
2. Deformation Theory	12
2.1. Covers of E of low gonality	19
3. Hyperplane Sections of $\tilde{V}_{Ne+df,g}^\circ(a, b)$	22
3.1. Nodal reduction	23
3.2. The genus bound	24
3.3. The dimension bound	28
3.4. Proof of Theorems 1.18 and 1.20	31
3.5. The general case	32
4. Monodromy Groups of Covers	36
5. Irreducibility of the Hurwitz Space $\mathcal{H}_{N,g}^0(E)$	41
5.1. Our strategy	41
5.2. The reduction argument	43
5.3. The degeneration argument	53
5.4. Proof of Lemma 5.14	54
5.5. Varying the moduli of E	55
5.6. Further remarks on the higher genus case	56
6. Components of the Severi Variety $V_{Ne+df,g}$	57
Appendix A. Deformation Theory	57
References	60

1. INTRODUCTION

We will study the geometry of two variants of classical parameter spaces. The first is the Hurwitz space $\mathcal{H}_{d,g}(C)$ of genus g , degree d simply branched covers of a fixed curve C . The classical case is when C is the projective line \mathbb{P}^1 , and goes as far back as to Clebsch in 1873. Much less is known in the case C has positive genus. In 1987, Gabai and Kazez [GK87] described the irreducible components of $\mathcal{H}_{d,g}(C)$ for arbitrary smooth curve C , through topological methods. One of our main results is an algebraic proof of their description in the case $C = E$ is a genus one curve.

The second space we consider is the Severi variety $V_{\mathcal{L},g}$ of geometric genus g integral curves of class \mathcal{L} on a smooth surface S . The classical case is when $S = \mathbb{P}^2$, and goes back to Severi himself, who gave an incomplete argument for its irreducibility. The gap was filled only decades later by Harris [Har86], and the techniques introduced there led to a series of discoveries in both qualitative (as in irreducibility statements) and quantitative (as in curve counting) aspects of the geometry of Severi varieties on various surfaces S . The method has been particularly successful in the case the surface S is rational. Here we will apply the same techniques, but in the case where $S = E \times \mathbb{P}^1$, for a connected smooth genus one curve E . Most of the results go through, but with many new interesting phenomena. In particular, we will show that the generic integral curve in $E \times \mathbb{P}^1$ is nodal, that the Severi variety has the expected dimension and we will describe the components of generic hyperplane sections in terms of *generalized Severi varieties*, much like in Caporaso–Harris [CH98]. We will also describe the irreducible components of the Severi variety in a restricted range of genus g and class \mathcal{L} .

Perhaps the main new feature of these two variants, the Hurwitz space $\mathcal{H}_{d,g}(E)$ and the Severi variety of curves in $E \times \mathbb{P}^1$, is that the study of one aids in the study of the other. The description of the components of hyperplane section of the Severi variety is the key input to determine the components of Hurwitz space. Conversely, we can leverage the knowledge of the components of Hurwitz space to determine the components of the Severi variety, at least in some range.

We start with a quick but self-contained introduction to the guiding questions and to a few known results in the study of both Severi varieties and Hurwitz spaces. This exposition is meant only as a motivation for the results and techniques that will occupy the main text, and it is by no means a complete survey of either field.

1.1. Hurwitz spaces. For a fixed smooth curve C , we define the Hurwitz space $\mathcal{H}_{d,g}(C)$ as the scheme parametrizing degree d covers (that is, finite flat maps) $f: X \rightarrow C$ that are simply branched, and whose source X is a smooth connected curve of genus g .

There is a branch morphism $\text{br}: \mathcal{H}_{d,g}(C) \rightarrow (C^b - \Delta)/S_b$ sending a cover $X \rightarrow C$ to its set of $b = 2g - 2 - d(2g(C) - 2)$ branch points. By the Riemann existence theorem, this is an étale map. In particular $\mathcal{H}_{d,g}(C)$ is smooth of dimension b .¹

¹Historically, Hurwitz used the Riemann existence theorem to put a complex manifold structure on $\mathcal{H}_{d,g}(\mathbb{P}^1)$, which was studied previously by Clebsch as a parameter space for topological covers of the sphere. For a modern treatment of the construction of the Hurwitz space, and in particular how to endow it with a scheme structure, see Fulton's [Ful69].

Next, we ask what are the connected components of $\mathcal{H}_{d,g}(C)$ —that is, which covers can be deformed into each other. As the Hurwitz space is smooth, every connected component is also irreducible. The famous theorem of Clebsch [Cle73] and Lüroth answers this question when C is rational.

Theorem 1.1 (Clebsch, Lüroth). *The Hurwitz space $\mathcal{H}_{d,g}(\mathbb{P}^1)$ is connected.*

As the forgetful map $\mathcal{H}_{d,g}(\mathbb{P}^1) \rightarrow \mathcal{M}_g$ is dominant for large d , this implies that \mathcal{M}_g is irreducible! This argument, first indicated by Klein, and then refined by Hurwitz and Severi, was the first known proof of the irreducibility of \mathcal{M}_g .

For higher genus targets, the Hurwitz space $\mathcal{H}_{d,g}(C)$ is usually not irreducible. For example, for each cover $X \rightarrow C$, consider the maximal unramified subcover $\tilde{C} \rightarrow C$ it factors through as $X \rightarrow \tilde{C} \rightarrow C$. The space of unramified covers is discrete (since the dimension of the Hurwitz space is the number of branch points). Therefore, to each cover point in $\mathcal{H}_{d,g}(C)$, we associated a discrete invariant, which will separate connected components.

Definition 1.2. We call a cover $f: X \rightarrow C$ *primitive* if it does not factor as $X \rightarrow \tilde{C} \rightarrow C$, with $\tilde{C} \rightarrow C$ a non-trivial unramified map. Let $\mathcal{H}_{d,g}^0(C)$ be the open and closed subscheme of $\mathcal{H}_{d,g}(C)$ parametrizing primitive covers.

We can reexpress this discrete invariant in a more topological vein. Over the complex numbers, for each cover $f: X \rightarrow C$, pick base points and consider the image of the fundamental group $f_*\pi_1(X, \star) \subset \pi_1(C, \star)$. Its conjugacy class is a discrete invariant of the cover $X \rightarrow C$. By the theory of covering spaces, the conjugacy class carries the same information as the maximal unramified subcover $\tilde{C} \rightarrow C$. In particular, being primitive is equivalent to having a surjective pushforward map on π_1 . For a simply branched cover, being primitive is also equivalent to having full monodromy—see Proposition 4.1.

Note that any cover $f: X \rightarrow C$ can be factored as $X \rightarrow \tilde{C} \rightarrow C$ where $X \rightarrow \tilde{C}$ is primitive, and $\tilde{C} \rightarrow C$ is unramified. Hence, if we understand the geometry of each $\mathcal{H}_{d,g}^0(C)$ for every d and C , we can recompose the whole Hurwitz space $\mathcal{H}_{d,g}(C)$ —it is just a disjoint union of $\mathcal{H}_{\tilde{d},g}^0(\tilde{C})$ over unramified maps $\tilde{C} \rightarrow C$ of degree d/\tilde{d} .

The main theorem of [GK87] is that this is the only discrete invariant.

Theorem 1.3 (Gabai–Kazez). *The Hurwitz space $\mathcal{H}_{d,g}^0(C)$ of primitive covers is connected (and hence irreducible) for any smooth curve C and $g > g(C)$.*

The traditional approach to irreducibility of Hurwitz space is group theoretic. Riemann existence theorem describes the fibers of the branch morphism in terms of the monodromy permutations associated to the branch points and the generators of $\pi_1(C)$. To establish connectedness of Hurwitz space, we need to understand how these permutations change as we move the branch points around, and to show that under a suitable sequence of moves, we can get from any permutation data to another. This is the approach in Clebsch’s original proof of irreducibility of $\mathcal{H}_{d,g}(\mathbb{P}^1)$. This method can be extended to covers of higher genus curves, as long as there are at least $2d - 2$ branch points. This was first discovered in Hamilton’s unpublished thesis, and rediscovered by Bernstein–Edmonds and Graber–Harris–Starr [GHS02].

A weaker version of Theorem 1.3 was conjectured by Bernstein and Edmonds in [BE84] as the *uniqueness conjecture*. There, besides moving the branch points around, they also allowed performing Dehn twists on the target curve. They got more mileage out of that, but to establish the full range of Theorem 1.3 some new key constructions had to be introduced by Gabai and Kazez in [GK87]. For a flavor of their techniques, see Lemma 5.21.

The language and methods in [GK87] are very much topological, which may have created somewhat of a gap between the algebraic geometry and topology literature. As a matter of fact, I have not seen Theorem 1.3 stated in the form above in the algebraic geometry literature, even though there has been a recent spur of interest in these Hurwitz spaces of covers of higher genus curves, as in for example [GHS02, Kan04, Kan05a, Kan05b, Che10, Kan13].

Now remains the question of whether algebraic methods can be used to establish Theorem 1.3. We highlight two results in this direction. In [Kan13], Kanev deals with the case where the degree of the cover is at most five (with some constraints on the number of branch points), using Casnati–Ekedahl–Miranda type of constructions. In [Kan03], Kani gives a beautiful proof that Theorem 1.3 holds when the genus of the target is one, and the genus of the source is two. As a matter of fact, he shows that $\mathcal{H}_{d,2}^0(E)$ is isomorphic to $E \times X(d)$, where $X(d)$ is the modular curve.

Here we will give an algebraic proof of Theorem 1.3 when the genus of the target C is one.

Theorem 1.4. *For a smooth genus one curve E and $g > 1$, the Hurwitz space $\mathcal{H}_{d,g}^0(E)$ of primitive covers is irreducible.*

Going back to the entire Hurwitz space $\mathcal{H}_{d,g}(E)$, we get the following corollary.

Corollary 1.5. *The components of $\mathcal{H}_{d,g}(E)$ are in bijection with the set of isogenies $\tilde{E} \rightarrow E$ of degree dividing d , but not equal to d . In particular, there are $\sum_{d \neq \tilde{d} | d} \sigma(\tilde{d})$ components, where $\sigma(n)$ is the sum of the positive divisors of n .*

An interesting—and perhaps more natural—variant is to let the target curve E vary as well.

Definition 1.6. Let $\mathcal{H}_{d,g,h}$ be the Hurwitz space parametrizing degree d simply branched covers from a pointed genus g to a genus h curve. More formally, a map $T \rightarrow \mathcal{H}_{d,g,h}$ corresponds to a smooth family of 1-pointed genus h curves $(\pi : \mathcal{D} \rightarrow T, \sigma : T \rightarrow \mathcal{D})$, a finite flat map $f : \mathcal{C} \rightarrow \mathcal{D}$ such that $f_t : C_t \rightarrow D_t$ is simply branched for every $t \in T$, and a section $\tau : T \rightarrow \mathcal{C}$ such that $\tau \circ f = \sigma$. Let $\mathcal{H}_{d,g,h}^0 \subset \mathcal{H}_{d,g,h}$ be the subscheme parametrizing primitive covers.

Remark 1.7. The extra data of the marked point is irrelevant for irreducibility questions, since the forgetful map to the unpointed variant will have irreducible one-dimensional fibers.

There is a map from the Hurwitz space $\mathcal{H}_{d,g,1}^0$ to $\mathcal{M}_{1,1}$. The fiber over (E, p) is the universal curve over $\mathcal{H}_{d,g}^0(E)$. Theorem 1.4 says the fibers are all irreducible of the same dimension, which implies the following corollary.

Corollary 1.8. *The Hurwitz space $\mathcal{H}_{d,g,1}^0$ is irreducible.*

Remark 1.9. In topological language, the irreducibility of $\mathcal{H}_{d,g,h}^0$ is equivalent to the uniqueness conjecture of [BE84]. Bernstein and Edmonds do establish it in the case $h = 1$, corresponding to the corollary above, by a careful combinatorial study of the monodromy data.

To determine the components of $\mathcal{H}_{d,g,1}$, the key discrete invariant of a cover $f : C \rightarrow E$. One of them is the isomorphism type of the cokernel $C(f)$ of the push-forward map f_* on π_1 (or H_1). It is an abelian group generated by at most two elements. The only possibilities are $C_{d_1} \oplus C_{d_2}$, with $d_1 \mid d_2$ and C_n is the cyclic group of order n .

We can re-encode this invariant in a more algebraic way. Let \tilde{d} be the degree of the maximal isogeny $\tilde{E} \rightarrow E$ that f factors through. These isogenies in general will get swapped by the monodromy of varying the target E , but there is a special family of isogenies that are fixed by the monodromy: the multiplication by m maps. We can remember as well the maximal m for which $C \rightarrow E$ factors through the multiplication by m map. We constructed a map that sends a cover $f : C \rightarrow E$ to a pair (\tilde{d}, m) . This carries the same data as the isomorphism type of the cokernel group $C(f)$ discussed above. Indeed,

$$C(f) = C_m \oplus C_{\frac{\tilde{d}}{m}}$$

In Section 5.5, we will prove the following.

Corollary 1.10. *The components of $\mathcal{H}_{d,g,1}$ are separated by the isomorphism type of $C(f)$. In particular, there is a bijection between the components of $\mathcal{H}_{d,g,1}$ and pairs (\tilde{d}, m) where $m^2 \mid \tilde{d} \mid d$.*

This answers one of the questions raised in [Che10].

Our argument for Theorem 1.4 will be analogous to Fulton’s proof of the irreducibility of \mathcal{M}_g given in the appendix of [HM82]. For more details and motivation for both our argument and Fulton’s proof, please turn to Section 5.1. But, in a nutshell, the proof of Theorem 1.4 goes like this. First consider the compactification of $\mathcal{H}_{d,g}(E)$ in the Kontsevich space $\overline{\mathcal{M}}_g(E, d)$ of stable maps $X \rightarrow E$ of degree d . Let us restrict ourselves to the locus where the map $X \rightarrow E$ has no contracted components—call this $\tilde{\mathcal{H}}_{d,g}(E)$.

Using induction on d and g , we will show that only one irreducible component of $\tilde{\mathcal{H}}_{d,g}^0(E)$ may contain covers with singular source. This amounts to identifying the boundary divisors, and use induction to show that each of them is irreducible, and that their union is connected. Since $\tilde{\mathcal{H}}_{d,g}^0(E)$ is smooth, only one component may contain a point in these boundary divisors.

To finish it off, we want to show that every component of $\tilde{\mathcal{H}}_{d,g}(E)$ does contain a cover with singular source. This is surprisingly hard to prove directly, and here is where our study of Severi varieties on $E \times \mathbb{P}^1$ enters.

1.2. Severi varieties. Given a line bundle \mathcal{L} on a surface S , we define the *Severi variety* $V_{\mathcal{L},g} \subset |\mathcal{L}|$ as the locus of integral curves of geometric genus g . The guiding questions to answer are the following.

- What is the dimension of the Severi variety?

- Do the generic points correspond to nodal curves?
- What are its irreducible components?
- What is its degree?

There has been a lot of work in these problems, and variants of them. Historically the problem originated by Severi's study of nodal planar curves, and nowadays we can answer these questions for most rational surfaces S and line bundles \mathcal{L} . There have been impressive results in the case S is a K3 surface, where the questions are much more delicate, and even some results in the much harder case where the surface S is of general type. For a general introduction, we refer the reader to Joachim Koch notes' of Ciro Ciliberto's lectures [Cil99].

We will focus on the technique introduced by Harris in [Har86], and perfected by Caporaso–Harris, Vakil, Tyomkin, and recently Shoval–Shustin [CH98, Vak00, Tyo07, SS13]. The originating question of this project was how much of these techniques extend over to non-rational surfaces, such as $E \times \mathbb{P}^1$.

Before we explain the technique in the context of $E \times \mathbb{P}^1$, let us deal with the fact its Picard group is not discrete, as opposed to the rational surface case. The Picard group of $E \times \mathbb{P}^1$ is generated by the pullback of $\mathcal{O}_{\mathbb{P}^1}(1)$ (which we will denote by e) and pullbacks from E . Therefore, it is one dimensional, just as $\text{Pic}(E)$.

This potentially could cause minor issues, since fixing the numerical class of a divisor does not pin it down. We avoid this issue entirely by noting that the automorphism group of E acts transitively on the set of divisors of same (non-zero) degree. Hence, the choice of line bundle is irrelevant to the questions above.

Based on this observation, we will denote $V_{df+Ne,g} \subset |\mathcal{L}|$ as the Severi variety of integral genus g curves of class \mathcal{L} , for some divisor \mathcal{L} with homology class $df + Ne$. When needed, we will be more explicit about \mathcal{L} and write $V_{\mathcal{L},g}$.

To answer the first two questions (the dimension and the singularities of the generic point), we will use the now standard deformation theory approach introduced in [CH98]. We conclude that the Severi variety $V_{df+Ne,g}$ has the expected dimension $2d + g - 1$ (with a single exception), and that the generic point does correspond to a nodal curve. We will do this in Section 2. The exception corresponds to the case $d = 0, N = g = 1$: when the Severi variety parametrizes the fibers of the projection $E \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$. Here the dimension is one, while the expected dimension works out to zero.

Note that, surprisingly, the expected dimension does not depend on N . This will have many interesting effects.

Our approach to the degree and irreducibility questions is directly analogous to Caporaso–Harris as well. We will fix an elliptic fiber $E_0 \subset E \times \mathbb{P}^1$ of the projection to \mathbb{P}^1 . For a generic point $p \in E_0$, consider the hyperplane $H_p \subset |\mathcal{L}|$ of curves passing through p . We want to describe the intersection

$$(1) \quad V_{df+Ne,g} \cap H_{p_1} \cap H_{p_2} \cap \dots \cap H_{p_a}$$

for generic points $p_1, \dots, p_a \in E_0$. We do this one hyperplane at a time. Eventually, the elliptic fiber E_0 is forced to split off, and most of our work goes into describing the residual curve.

To understand this intersection, we need a better compactification of the Severi variety $V_{df+Ne,g}$. We could do this by simply taking the closure $\overline{V_{df+Ne,g}} \subset |\mathcal{L}|$.

Unfortunately this closure is not normal, so it will be more convenient for us to work with its normalization $\tilde{V}_{df+Ne,g}$.

Here is a component we expect to see in the intersection (1).

Definition 1.11. For $p_1, \dots, p_a \in E_0$, and a line bundle $L \in \text{Pic}^b E_0$ such that $a+b = d$, we define a *generalized Severi variety* $\tilde{V}_{\mathcal{L},g}(a,b)[p_1, \dots, p_a, L]$ as the normalization of the closure of the locus of integral curves $C \subset E \times \mathbb{P}^1$ of geometric genus g and class \mathcal{L} , such that the intersection of C with E_0 is transverse, composed of $a+b = d$ points, a of which are the points p_1, \dots, p_a , and the sum of the remaining b points has class L . When no confusion arises, we will write only $\tilde{V}_{df+Ne,g}(a,b)$, where the corresponding homology class of \mathcal{L} is $df + Ne$, and leave the data of \mathcal{L} , L and the points p_i implicit.

We will show that when $b > 0$ the Severi variety $\tilde{V}_{df+Ne,g}(a,b)$ has expected dimension and that the generic point corresponds to a nodal curve. However, these are not the only components we will see in the intersections with H_{p_i} . Whenever the fiber E_0 is split off generically, the residual curve might not be transverse to E_0 . We will need some notation to keep track of the tangency profile of the residual curve with E_0 . Here our notation will differ slightly from the standard notation in Caporaso–Harris.

Definition 1.12. A *(tangency) profile* α is an ordered sequence of positive integers $(\alpha_1, \alpha_2, \dots, \alpha_k)$. A sum of profiles $\alpha^1 + \alpha^2$ is the profile obtained by concatenating the sequences. The *size* of the profile $|\alpha|$ is k , the number of entries in the sequence. The *multiplicity* of a profile $m(\alpha)$ is the sum $\sum_{i=1}^k \alpha_i$.

A subprofile $\alpha' \subset \alpha$ is a subsequence of $(\alpha_1, \dots, \alpha_k)$. The complement $\alpha - \alpha'$ is the profile given by the complementary subsequence.

For consistency, we will say that an integer k corresponds to the profile $1^k = (1, 1, \dots, 1)$ with k entries.

We will use tangency profiles to book-keep tangencies in the most naive way—having a tangency profile α means there is a point with tangency of order α_1 , another with order α_2 , and so on².

Definition 1.13. For points p_1, \dots, p_a and a tangency profile β such that $a + m(\beta) = d$, and $L \in \text{Pic}^{m(\beta)} E_0$, define $\tilde{V}_{df+Ne,g}(a,\beta)[p_1, \dots, p_a, L]$ as the normalization of the closure of the locus of integral curves $C \subset E \times \mathbb{P}^1$ of geometric genus g , class \mathcal{L} , whose intersection with E_0 is

$$E_0 \cap C = \sum_{i=1}^a p_i + \sum_{j=1}^{|\beta|} \beta_j q_j$$

and the class of $\sum_{j=1}^{|\beta|} \beta_j q_j$ is L .

²Caporaso–Harris in [CH98] re-encoded these tangency conditions in a different sequence $\bar{\alpha}$, such that $\bar{\alpha}_k$ is the number of α_i that are equal to k . Their convention definitely has computational advantages, but it would also make the notation more loaded later on. As we will not derive any recursive degree formulas in this article, we keep the more naive notation for the sake of notational simplicity.

Remark 1.14. We have kept this line bundle L in our notation for the Severi variety $V_{df+Ne,g}(\alpha, \beta)[p_1, \dots, p_a, L]$, but of course we can recover it from the class \mathcal{L} of the curve in $E \times \mathbb{P}^1$ and the points p_i , as we have

$$\mathcal{L}|_{E_0} = \sum_{i=1}^a p_i + L$$

This choice will be justified in Definition 1.17.

Note that under the convention that an integer b corresponds to the profile 1^b , Definitions 1.11 and 1.13 agree.

Still, these are not all the possible varieties parametrizing the residual curves. There are two more phenomena that can happen. First, even though $V_{df+Ne,g}$ parametrizes integral curves, there is no reason to expect that once a fiber E_0 splits off, the residual curve is still irreducible.

We could define $\tilde{V}_{df+Ne,g}^{\text{reduced}}(\alpha, \beta)$ as in definition 1.13, but imposing that the curves are only *reduced*, instead of integral. Now g corresponds to the arithmetic genus of the normalization. This should keep track of the reducible residual curves as well.

One would hope that the dimension count will still work for $\tilde{V}_{df+Ne,g}^{\text{reduced}}(\alpha, \beta)$. However, it does not, because of the exception mentioned earlier: the fibers of $E \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$ move in a one-dimensional family, while the dimension count predicts a zero dimensional space. Let us introduce a notation for the components of $\tilde{V}_{df+Ne,g}^{\text{reduced}}(\alpha, \beta)$ which do not contain generically any of these fibers.

Definition 1.15. Let $\tilde{V}_{df+Ne,g}^{\circ}(\alpha, \beta) \subset \tilde{V}_{df+Ne,g}^{\text{reduced}}(\alpha, \beta)$ be the closure of the components whose generic point does not contain a fiber of the projection $E \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$.

Remark 1.16. Note that $\tilde{V}_{df+Ne,g}^{\circ}(\alpha, \beta)$ parametrizes curves whose (possibly disconnected) normalizations have arithmetic genus g . If the normalization is disconnected, there may be even components with genus bigger than g , as long as there are rational components to balance off. Of course, there could be at most d of those, so no components of arbitrarily high genus ever arise in $\tilde{V}_{df+Ne,g}^{\circ}(\alpha, \beta)$.

Still, it makes sense, and will be needed, to consider $\tilde{V}_{df+Ne,g}^{\circ}(\alpha, \beta)$ for *negative* values of g . The curves thus parametrized always contain some number of fibers from $E \times \mathbb{P}^1 \rightarrow E$.

There is a second possible phenomenon, as seen in the following example. Say we are looking at the intersection of $\tilde{V}_{Ne+df,g}(\alpha, b)[p_1, \dots, p_a, L]$ with a hyperplane H_p . Let W be a component of the intersection. It could be equal to a component of $\tilde{V}_{Ne+df,g}(\alpha+1, b-1)$, meaning that one point which was “free to move” became fixed at p . Or it could be that the fiber E_0 was split off generically with multiplicity m . Let X be the residual curve. We will completely characterize it soon, but for now let us consider just the case where the limit of all the b points land in X , meaning, the intersection of X with E_0 contains the limit of the b points. Now, the linear class of the sum of these b points was fixed and equal to L —it still has to be so in the limit. Moreover, the $d-b$ remaining points of intersection of X with E_0 might not be all fixed. In this case, we will need to keep *two* tangency profiles around to keep track of the intersection of X with E_0 —one for the limit of the b points in the

generic fiber, whose class sums up to L , and another for any new tangencies that might occur. We accommodate for this with the following definition.

Definition 1.17. For points p_1, \dots, p_a , and tangency profiles β^1, β^2 such that

$$a + m(\beta^1) + m(\beta^2) = d$$

and line bundles $L_k \in \text{Pic}^{m(\beta^k)} E_0$, define $\tilde{V}_{df+Ne,g}^\circ(a, \beta^1, \beta^2)[p_1, \dots, p_a, L_1, L_2]$ as the normalization of the closure of the locus of reduced curves $C \subset E \times \mathbb{P}^1$ (not containing any fiber of $E \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$) of geometric genus g , class \mathcal{L} , whose intersection with E_0 is

$$E_0 \cap C = \sum_{i=1}^a p_i + \sum_{j=1}^{|\beta^1|} \beta_j^1 q_j^1 + \sum_{j=1}^{|\beta^2|} \beta_j^2 q_j^2$$

and the class of $\sum_{j=1}^{|\beta^k|} \beta_j^k q_j^k$ is L_k , for $k = 1, 2$.

We are now ready to state a complete characterization of the possible components of $H_p \cap \tilde{V}_{df+Ne,g}^\circ(a, b)$. The description is easier to state if we look at the intersection with $\tilde{V}_{df+Ne,g}^\circ(a, b)$, so let us do so.

Theorem 1.18. For $g \geq 2$, and p be a generic point in E_0 , let

$$W \subset H_p \cap \tilde{V}_{df+Ne,g}^\circ(a, b)[p_1, \dots, p_a, L]$$

be an irreducible component of the intersection. Then W falls into one of the following three cases.

If the curve E_0 does not split off generically in W , then W is a component of

$$(2) \quad \tilde{V}_{df+Ne,g}^\circ(a+1, b-1)[p_1, \dots, p_a, p, L-p]$$

This case can only happen if $b \geq 2$.

Otherwise, the curve E_0 splits off generically, say with multiplicity m . We can view W as parameter space for the residual curve, whose class is $\mathcal{L} - mE_0$. Then there is a subsequence (i_1, \dots, i_a) of $(1, 2, \dots, a)$ and a tangency profile $\tau \neq (1)$ such that if we let $\bar{g} = g - |\tau|$, then either:

- W is a component of

$$(3) \quad \tilde{V}_{df+(N-m)e,\bar{g}}^\circ(\bar{a}, b-1+\tau)[p_{i_1}, \dots, p_{i_a}, \bar{L}]$$

where $\bar{a} = d - (b - 1 + m(\tau)) \geq 0$ and $\bar{L} = \mathcal{L}|_{E_0} - \sum_{k=1}^{\bar{a}} p_{i_k}$,

- or W is a component of

$$(4) \quad \tilde{V}_{df+(N-m)e,\bar{g}}^\circ(\bar{a}, b, \tau)[p_{i_1}, \dots, p_{i_a}, L, \bar{L}]$$

where $\bar{a} = d - (b + m(\tau)) \geq 0$ and $\bar{L} = \mathcal{L}|_{E_0} - L - \sum_{k=1}^{\bar{a}} p_{i_k}$.

Remark 1.19. Theorem 1.18 gives only a set-theoretic description of the boundary. To have a complete scheme theoretic characterization, we would need to compute the multiplicity in which each component appears. This computation is planned to appear in an upcoming follow-up article.

The main new feature here, compared with results in [CH98, Vak00, SS13], is that E_0 may be generically split off with multiplicity $m > 1$ —this will correspond in the stable reduction to getting an m -fold unramified cover dominating E_0 .

The next question is: what happens when we intersect $\tilde{V}_{df+(N-m)e,g'}^\circ(a', b, \tau)$ with a hyperplane section H_p ? We will answer this question in Theorem 3.13, and furthermore we can describe all the components that ever arise in the process of intersecting with hyperplane sections H_p . To do so, however, we need to extend even further our Definition 1.17 of Severi varieties $\tilde{V}_{df+Ne,g}^\circ(a, \beta^1, \beta^2)$ to allow for $a = 1^a$ to be an arbitrary tangency profile, and to allow arbitrarily many β^k . The statement of the result becomes notationally more loaded, and we leave it to Section 3.5.

The good news is that the proof of Theorem 1.18 already contains all the interesting challenges, and adapting it to the general case of Theorem 3.13 is just a matter of additional bookkeeping. We will then first prove the simpler Theorem 1.18 in Section 3, and then translate its proof to the stronger Theorem 3.13 in Section 3.5.

The missing ingredient to get a recursive formula for the degree of $\tilde{V}_{df+Ne,g}$ is to find out if all these components appear, and with what multiplicity. Thinking in terms of a recursive formula for the Severi degrees, Theorem 3.13 tells us what terms would appear in the formula, but does not tell us with what coefficients. The multiplicity computation is planned to appear in an upcoming follow-up article. As a starting point for the recursion, we would need to know the degrees when $g = 1$, which have been already calculated by Cooper–Pandharipande in [CP12].

The technique used to prove Theorem 1.18 was introduced in [CH98]. We will exploit the tension between two different points of views on Severi varieties: embedded curves versus parametrized curves. From the embedded curves perspective, we know the dimension of the Severi variety $V_{df+Ne,g}^\circ(a, b)$, which gives us control on how many conditions we may impose on the curve corresponding to the general point of the divisor W . From the parametric point of view, we can study the general point of W by taking a general arc in $V_{df+Ne,g}^\circ(a, b)$ approaching it, and applying stable reduction to the family of normalization maps. The central fiber is replaced by a nodal curve, which allows us to concretely understand the combinatorial restrictions of having arithmetic genus g . It turns out that the number of constraints we discover from the parametric point of view exactly matches the number of conditions the embedded point of view allows. This precise balance pins down exactly what are the possible limit components listed in Theorem 3.13.

It is worth mentioning that although our exposition uses both points of view, it is definitely possible to write it all in terms of the parametric one. This is done, for example, in [Vak00]. The proofs usually get shorter, but at a slight cost of some additional technical details. Our compromise of using both points of view is just a matter of taste—we want to make the implicit tension between the two points of view visible already at the level of language.

There is a rational map $\tilde{V}_{df+Ne,g}^\circ(a, b) \dashrightarrow \overline{\mathcal{M}}_g(E, d)$, sending integral curves $C \subset E \times \mathbb{P}^1$ to the composition of the normalization $C^\vee \rightarrow C$ with the projection $C \subset E \times \mathbb{P}^1 \rightarrow E$. The rational map can be completed in codimension 2, and hence it is well defined on the generic points of the components of the hyperplane section

$H_p \cap \tilde{V}_{df+Ne,g}(a, b)$. Out of the proof of Theorem 1.18, we also get a description of this map, as follows.

Theorem 1.20. *Using the notation of Theorem 1.18, the image of the general point $[C] \in W$ under the rational map to $\overline{\mathcal{M}}_g(E, d)$ is as follows.*

- *If W is a component of $V_{df+Ne,g}(a+1, b-1)$, then the general point corresponds to an integral curve $C \subset E \times \mathbb{P}^1$, and its image in $\overline{\mathcal{M}}_g(E, d)$ is the composition of the normalization map $C^\vee \rightarrow C$ with the projection $C \subset E \times \mathbb{P}^1 \rightarrow E$.*
- *Otherwise the curve E_0 splits off generically with multiplicity m , and the residual curve is a generic point $[R]$ of a component of $\tilde{V}_{df+(N-m)e,g}^\circ(\bar{a}, b-1+\tau)$ or of $\tilde{V}_{df+(N-m)e,g}^\circ(\bar{a}, b, \tau)$. Then its image in $\overline{\mathcal{M}}_g(E, d)$ is a stable map $X \rightarrow E$, where X is a connected nodal curve such that:*
 - *the normalization of X is the disjoint union of \tilde{E} and the normalization of R , where $\tilde{E} \rightarrow E_0$ is a degree m unramified cover (which may be disconnected),*
 - *and X has $|\tau|$ nodes, all of them connecting components of \tilde{E} to components of the normalization of R .*

Theorem 1.20 is the key ingredient in our proof of Theorem 1.4, the irreducibility of the Hurwitz space of primitive covers of elliptic curves. As briefly discussed in Section 1.1, we can reduce the irreducibility statement to showing that every connected component of $\tilde{\mathcal{H}}_{d,g}(E)$ contains a cover with singular source. And Theorem 1.20 produces this cover for us!

After using the degeneration coming from the Severi variety world to prove Theorem 1.4, we come back and use the irreducibility of $\mathcal{H}_{d,g}^0(E)$ to prove a result about Severi varieties: we determine the irreducible components of $V_{Ne+df,g}$ when $d \geq 2g - 1$. Again, their description is very much like in Section 1.1—the unique discrete invariant will be the maximal unramified map through which the map from the normalization of the embedded curve to E factors. For more details, see Section 6.

1.3. Guide to readers. The sections are organized to be as independent of each other as possible. The reader interested in a particular section should be able to read it directly, as long as one is willing to assume a couple of statements proven in other sections. Here is a short summary of the sections' content and dependencies.

- In Section 2, we start our study of Severi varieties by determining its dimension, and showing that its generic point corresponds to a nodal curve. All results are based on a now standard theorem in deformation theory, of which many versions already exist in the literature. For completeness, we prove the precise formulation we use (Theorem 2.1) in Appendix A.
- In Section 3, we analyze the components of a hyperplane section of the Severi variety, and prove Theorems 1.18 and 1.20 and some generalizations. The argument depends on the dimension bounds Corollary 2.9 and Proposition 2.10, both proven in Section 2.
- In Section 4, we study the monodromy group of both primitive and non-primitive simply branched covers. This discussion does not depend on the previous sections.

- In Section 5, we prove Theorem 1.4, that is, the Hurwitz space of primitive covers of an elliptic curve is irreducible. Most of the argument is independent of the previous sections, but at two key moments we invoke Lemma 4.7 (proven in Section 4) and Theorem 1.20 (proven in Section 3).
- Finally, in Section 6, we use Theorem 1.4 to describe the irreducible components of the Severi variety $V_{df+N_{E,g}}$ whenever $d \geq 2g - 1$.

The results are valid over algebraically closed fields of characteristic zero. For simplicity though, we will work over the complex numbers throughout this article. The only places this choice shows up are in Section 4 and Section 5.4. In the former, the use of complex numbers may be bypassed by the machinery of étale fundamental groups, and in the latter by identifying isogenies with subgroups of the $E[N]$ instead of sublattices in the complex plane.

With the alterations above, the arguments of sections 4 and 5 should go through in characteristic p large enough compared to the degree of the cover. However, the key theorem 2.1—on which the whole Severi variety discussion is based on—is only known to be valid in characteristic zero. This is the same challenge that shows up, for example, in [Har86, Vak00, Tyo07].

For the reader interested in Hurwitz spaces in characteristic $p > 0$, probably adapting the results of Fulton's [Ful69] is a more likely route. There he defines the Hurwitz space of covers of \mathbb{P}^1 as a scheme over \mathbb{Z} , and deduces irreducibility in characteristic p from the result in characteristic zero.

1.4. Acknowledgments. I would like to thank Joe Harris, who welcomed me into the breathtaking world of algebraic geometry, and guided me throughout this entire journey. The original questions that initiated this project arose in conversations with Anand Patel, who has been a great collaborator and friend in the past few years. I thank him again for the thoughtful comments on an earlier draft, along side with Dawei Chen, Anand Deopurkar and Ravi Vakil. I am very grateful to Tom Graber, Jason Starr and Vassil Kanev for introducing me to the relevant topology literature. Finally, a special thanks to Carolina Yamate, who kindly drew the (several) beautiful pictures that give life to this manuscript.

2. DEFORMATION THEORY

We start our study of Severi varieties of curves in $E \times \mathbb{P}^1$ by determining its dimension, and describing the geometry of the generic curve. The techniques used are the same as in [CH98, Vak00, Tyo07]. The key result is the following.

Theorem 2.1. *Let Σ be a smooth projective surface, and fix a homology class $\tau \in H^2(\Sigma, \mathbb{C})$ and an integer g . Let $W \subset \mathcal{M}_g(\Sigma, \tau)$ be a non-empty subvariety of Kontsevich space of stable maps. Let $f: C \rightarrow \Sigma$ be a general point of W . Assume that f is birational onto its image.*

Let $D \subset \Sigma$ be a curve, and $\Omega \subset D$ a finite subset. Let $Q = f^{-1}(D - \Omega) \subset C$. Let $|Q|$ denote the cardinality of the set Q (that is, we are not counting multiplicity). Fix some integer b , such that $b \geq |Q|$, and let $\gamma = -(K_\Sigma + D)\tau + b$.

If $\gamma \geq 1$, then

$$\dim W \leq g - 1 + \gamma$$

Conversely, suppose that equality occurs. Then, the more positive γ , the better behaved the map f is:

- If $\gamma \geq 2$, then $|Q| = b$ and f is unramified away from Q .
- If $\gamma \geq 3$, then f is unramified everywhere, and the image $f(C)$ has at most ordinary multiple points, and is smooth along $f(Q)$.
- If $\gamma \geq 4$, then the image of f is nodal, and smooth along $f(Q)$.

Remark 2.2. Note that as $W \subset \mathcal{M}_g(\Sigma, \tau)$ and not its closure, the source C of $f: C \rightarrow \Sigma$ is smooth and irreducible. Later we will see what one can say about more general sources, in Proposition 2.8.

Remark 2.3. We will apply Theorem 2.1 only in the following setup. Fix a line bundle \mathcal{L} on Σ , and let $W \subset |\mathcal{L}|$ be a non-empty subvariety, and $f: C \rightarrow \Sigma$ the normalization of the general point of W . Hence, for our applications, we could have bypassed the language of Kontsevich spaces if we wanted to. However, the result as stated above is sharper, because it allows for the line bundle \mathcal{L} to vary continuously, which is important when the Picard group of Σ is not discrete.

Setting $\Omega = \emptyset$, and $b = D \cdot \tau$, the first part is explicitly stated as Corollary 2.4 in [CH98]. The other parts are also in [CH98], but implicit in the proof of Proposition 2.2. There are many similar results in the literature, such as Theorem 3.1 in [Vak00], and Theorem 1 and Lemma 2 in [Ty07]. Unfortunately, I could not locate a version of this result for the set up we are working on. However, it is easy to adapt the proofs in any of these sources to the exact statement above. For completeness, we do include a proof of Theorem 2.1, but leave it to Appendix A.

Let us apply the general Theorem 2.1 to our study of curves in $E \times \mathbb{P}^1$, with tangency conditions along a fiber $E_0 \subset E \times \mathbb{P}^1$. Here is a preliminary result.

Lemma 2.4. *The Severi variety $\tilde{V}_{df+N_e, g}(\alpha, 0)$ has dimension at least $d + g - 1$.*

Proof. We want to parametrize integral curves $C \subset E \times \mathbb{P}^1$ of divisor class \mathcal{L} , geometric genus g , and whose intersection with E_0 is fixed. We get our lower bound by a naive dimension count. Fixing the intersection with E_0 takes $h^0(E_0, \mathcal{L}|_{E_0}) = d$ conditions. And having geometric genus g imposes $\delta = p_a(\mathcal{L}) - g$ conditions. Indeed, if we knew our curve was nodal, then this count asks one condition for each node. We do not know the curve is nodal yet, but in any case we know that an isolated planar singularity of delta-invariant δ imposes δ conditions (in the language of [DH88], the *equigeneric locus* EG has codimension δ in the versal deformation space of the singularity).

In the worst case, all these conditions would be independent, and we get

$$\begin{aligned} \dim \tilde{V}_{df+N_e, g}(\alpha, 0) &\geq h^0(E \times \mathbb{P}^1, \mathcal{L}) - d - \delta \\ &= d(N + 1) - d - (Nd - d + 1 - g) = d + g - 1 \end{aligned}$$

as we wanted to show. \square

Lemma 2.5. *For $d \geq 1$, the dimension of $\tilde{V}_{df+N_e, g}(\alpha, 0)[p_1, \dots, p_k]$ is $d + g - 1$, and the generic point corresponds to a nodal curve.*

Proof. Lemma 2.4 tell us that the dimension of $\tilde{V}_{df+Ne,g}(\alpha, 0)$ is at least $d + g - 1$. We only need to show that it is also at most that. We will do that by applying Theorem 2.1. Set $\Sigma = E \times \mathbb{P}^1$, $D = E_0$, and $\Omega = \{p_1, \dots, p_k\}$.

Then $Q = f^{-1}(D - \Omega) = \emptyset$, and we may set $b = 0$. We get

$$\gamma = -(K_{E \times \mathbb{P}^1} + D) \cdot (df + Ne) + b = -(-2e + e)(df + Ne) + 0 = d \geq 1$$

Hence, by Theorem 2.1, we have

$$\dim \tilde{V}_{df+Ne,g}(\alpha, 0) \leq \gamma + g - 1 = d + g - 1$$

By Lemma 2.4, the opposite inequality holds as well, so we actually have equality.

Let us prove that the general point of $\tilde{V}_{df+Ne,g}(\alpha, 0)$ is nodal. By Theorem 2.1, as soon as $\gamma = d \geq 4$, we are good. Let us consider the remaining cases:

- For $d = 1$, the only class with integral curves is f , which only has smooth elements.
- For $d = 2$, Theorem 2.1 tell us it is enough to rule out tacnodes and triple points. Triple points can't occur because the map $C \rightarrow \mathbb{P}^1$ is two to one. We will rule out tacnodes by an ad hoc argument, in Section 2.1.1.
- For $d = 3$, we only have to rule out triple points, by Theorem 2.1. We will do this in Section 2.1.2.

□

Now we are ready to analyze the Severi variety $V_{Ne+df,g}(a, b)$.

Theorem 2.6. *The Severi variety $V_{Ne+df,g}(a, b)$ has dimension $d + g - 2 + b$, and the generic point corresponds to a nodal curve.*

Proof. Consider the map $V_{Ne+df,g}(a, b) \rightarrow \text{Sym}^{a+b} E$ sending a curve $C \subset E \times \mathbb{P}^1$ to its $a+b$ intersection points with E_0 . The fiber over (p_1, \dots, p_{a+b}) is $V_{df+Ne,g}(a+b, 0)$, which by Lemma 2.5 has dimension $d + g - 1$ and the generic point is nodal. We only have to show that the image of $V_{Ne+df,g}(a, b) \rightarrow \text{Sym}^{a+b} E$ has dimension $b - 1$. This is because a of the points are fixed, while the other b are constrained to be in a fixed linear series of degree b . Riemann–Roch tells us they vary in a $(b - 1)$ -dimensional series. □

Remark 2.7. One of the key difficulties in studying curves in $C \times \mathbb{P}^1$, for curves C of genus at least two, is the lack of an analogue of Theorem 2.6. Indeed, the techniques here stem out of Theorem 2.1, which works better the more negative the canonical bundle of the surface is.

In the discussion above, we were studying integral curves in $E \times \mathbb{P}^1$. In our application, however, we will not know a priori that the residual curves are integral, but we still want to have a bound on their dimension. In the parametric approach, integral curves correspond to maps $f: C \rightarrow E \times \mathbb{P}^1$ with smooth connected source and which are birational onto their image. Let us relax these conditions to just having a nodal source C , possibly disconnected, and with an arbitrary map $f: C \rightarrow E \times \mathbb{P}^1$. We ask for an upper bound on the deformations of the map f . More precisely, assume the following set up.

Hypothesis 1. Fix a line bundle \mathcal{L} on $E \times \mathbb{P}^1$, and a sequence of positive integers $\alpha_1, \dots, \alpha_n$. Let $\mathcal{C} \rightarrow V$ be a flat family of (possibly disconnected) semistable genus g curves, equipped with a map $\varphi: \mathcal{C} \rightarrow E \times \mathbb{P}^1$, and sections $\sigma_i: V \rightarrow \mathcal{C}$, such that

- For each $b \in V$, the pullback divisor $\varphi^{-1}(E_0)$ in \mathcal{C}_b is equal to

$$\varphi_b^{-1}(E_0) = \alpha_1 \sigma_1(b) + \alpha_2 \sigma_2(b) + \dots + \alpha_n \sigma_n(b)$$

- Class of $\varphi(\mathcal{C}_b) \subset E \times \mathbb{P}^1$ is \mathcal{L} .
- The induced map $V \rightarrow |\mathcal{L}|$ has finite fibers (equivalently, the image family $\varphi(\mathcal{C}_b) \subset E \times \mathbb{P}^1$ is nowhere isotrivial).

We want to bound the dimension of V , and to give a description of the equality case. For sake of preliminary discussion, let us fix completely the intersection with E_0 . That is, let us say that the sections $\sigma_i(V)$ are contracted under φ down to points p_i , for each $i = 1, \dots, n$.

We would hope that the bound of Lemma 2.5 would still hold. This is not the case. For example, consider letting $\mathcal{C} = \mathcal{C}' \sqcup \tilde{E}$ be a disjoint union, where $\mathcal{C}' \rightarrow E \times \mathbb{P}^1$ is the normalization of a point in $V_{(N-n)e+df,g}(\alpha, 0)$, and $\tilde{E} \rightarrow E \times \{t\} \subset E \times \mathbb{P}^1$ is a degree n isogeny landing in any fiber $E \times \{t\}$. Then the map $\mathcal{C}' \rightarrow E \times \mathbb{P}^1$ varies in a $d+g-1$ dimensional family, while $\tilde{E} \rightarrow E \times \mathbb{P}^1$ varies in one dimension (corresponding to choosing the target fiber $E \times \{t\}$), for a total of a $d+g$ dimensional family.

Fortunately, this is the only way to go beyond the $d+g-1$ bound. Let a *floating component* be a connected component of \mathcal{C}_b whose image under $\varphi(\mathcal{C}_b)$ has class multiple of e . We will show that if the family $\mathcal{C} \rightarrow V$ has no floating components, then the dimension of V is at most $d+g-1$. Let us first describe the equality cases.

As we are allowing disconnected sources \mathcal{C} , we could take an arbitrary disjoint union $\mathcal{C} = \mathcal{C}_1 \sqcup \mathcal{C}_2 \sqcup \dots \sqcup \mathcal{C}_k$, and let $\mathcal{C}_i \rightarrow E \times \mathbb{P}^1$ be a generic point of $V_{d_i f + N_i e, g_i}(\alpha^i, 0)$. We require that no \mathcal{C}_i is a floating component (i.e. we exclude the $d_i = 0$ case), and moreover that

$$\sum d_i = d, \quad \sum N_i = N, \quad \sum (g_i - 1) = g - 1, \quad \sum \alpha^i = \alpha$$

Then the map $\mathcal{C} \rightarrow E \times \mathbb{P}^1$ will vary in a

$$\sum d_i + g_i - 1 = d + g - 1$$

dimensional family as well.

We already have a notation to parametrize such curves. Recall definition 1.15: the Severi variety $\tilde{V}_{df+Ne,g}^\circ(\alpha, 0)$ is the normalization of the closure of the locus of reduced curves in $E \times \mathbb{P}^1$, of some specified class \mathcal{L} , some fixed intersection with E_0 , and whose normalization has arithmetic genus g , which contains no fibers of the projection $E \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$ (that is, no floating components). Compare to $\tilde{V}_{df+Ne,g}(\alpha, 0)$, which further imposes that the curves are irreducible. In general, when we want to allow disconnected sources, but avoid floating components, we will decorate our Severi varieties with a \circ superscript.

The variety $\tilde{V}_{df+Ne,g}^\circ(\alpha, 0)$ has dimension $d+g-1$, which we will soon prove to be maximal. One could hope that these are all equality cases. Unfortunately, there is one more circumstance under which equality is achieved. Take a curve




FIGURE 1. A curve with *elliptic tails*. We impose that each $E^i \rightarrow E \times \mathbb{P}^1$ is unramified.

$C \rightarrow E \times \mathbb{P}^1$ in $\tilde{V}_{df+(N-N')e, g-k}^\circ(\alpha, 0)$, choose k points $q_1, \dots, q_k \in C$, and attach isogenies $E^i \rightarrow E$ at each of these points, such that the sum of the degrees of the k isogenies is N' . Let

$$\bar{C} = C \cup_{q_1} E^1 \cup \dots \cup_{q_k} E^k \rightarrow E \times \mathbb{P}^1$$

be this new curve, with its map to $E \times \mathbb{P}^1$ (see fig. 1). Then it also varies in

$$k + \dim \tilde{V}_{df+(N-N')e, g-k}^\circ(\alpha, 0) = k + d + g - k - 1 = d + g - 1$$

dimensions. Call the curves E^i *elliptic tails*, and set the space of all such maps $\bar{C} \rightarrow E \times \mathbb{P}^1$ as $\tilde{V}_{df+Ne, g}^{\circ, \text{tails}}(\alpha, 0)$. Generally, we will use the superscript *tails* to denote that we allow elliptic tails in the source.

The components of $\tilde{V}_{df+Ne, g}^{\circ, \text{tails}}(\alpha, 0)$ will be all equality cases. The result is the following.

Proposition 2.8. *Assume Hypothesis 1 and moreover that for generic $b \in V$, the map $C_b \rightarrow E \times \mathbb{P}^1$ has no floating components, and that the sections $\sigma_i(V)$ get contracted by the map $\varphi: C \rightarrow E \times \mathbb{P}^1$ to points p_i , for all i . Then, $\dim V$ is at most $d + g - 1$.*

If equality occurs, then the image of V in $|\mathcal{L}|$ is dense in a component of the Severi variety $\tilde{V}_{df+Ne, g}^{\circ, \text{tails}}(\alpha, 0)[p_1, \dots, p_n]$.

The following proof follows Proposition 3.4 in Vakil's [Vak00] closely.

Proof. After passing to a dominant generically finite cover of V , we may distinguish the components of C_b . If generically there are contracted components, discard them. This only makes the bound tighter.

Let us replace C_b with its normalization \widetilde{C}_b . Again, this will only make the bound tighter. Let k be the number of floating components in \widetilde{C}_b , and let Y be the union of the other components. As the original curve C_b did not have any floating component, there were at least k nodes connecting them to other components of \widetilde{C}_b . Hence, the arithmetic genus of $p_a(Y)$ is at most $g - k$, with equality only if C_b is equal to the curve Y with k elliptic tails attached.

Each floating component of \widetilde{C}_b contributes to at most one dimension to V , because the map $V \rightarrow |\mathcal{L}|$ is generically finite, so the only moduli comes from picking the target fiber $E \times \{t\}$.

Let us focus on each component X of Y then. Let \bar{X} be the image of $\varphi: X \rightarrow E \times \mathbb{P}^1$, and $\tilde{X} \rightarrow \bar{X}$ its normalization. As X is smooth, the map $X \rightarrow \bar{X}$ factors through $\tilde{X} \rightarrow \bar{X}$. Moreover, since the fibers of $V \rightarrow |\mathcal{L}|$ are finite, a bound in the dimension in which $\tilde{X} \rightarrow \bar{X} \subset E \times \mathbb{P}^1$ varies also bounds the dimension of V . Let m be the degree of $X \rightarrow \tilde{X}$. Summing over the components of \widetilde{C}_b , and applying Lemma 2.5, we get

$$\begin{aligned}
\dim V &\leq \left(\sum_{X \text{ floating}} \dim \{X \rightarrow E \times \mathbb{P}^1\} \right) + \left(\sum_{X \text{ not floating}} \dim \{\tilde{X} \rightarrow E \times \mathbb{P}^1\} \right) \\
&\leq \left(\sum_{X \text{ floating}} 1 \right) + \left(\sum_{X \text{ not floating}} \frac{d_i}{m} + p_a(\tilde{X}) - 1 \right) \\
&\leq k + \left(\sum_{X \text{ not floating}} \frac{d_i}{m} + \frac{p_a(X) - 1}{m} \right) \\
&\leq k + \left(\sum_{X \text{ not floating}} d_i + p_a(X) - 1 \right) \\
&= k + d + p_a(\tilde{Y}) - 1 \\
&\leq k + d + (g - k) - 1 = d + g - 1
\end{aligned}$$

as we wanted to show.

If equality holds, then $m = 1$, Y is a generic point of $V_{df+N'e, g-k}^\circ(\alpha, 0)$ and C_b is a generic point in $\tilde{V}_{df+N'e, g}^{\circ, \text{tails}}(\alpha, 0)$. \square

In our application, the intersection with E_0 will not be fixed. However, we can immediately adapt Proposition 2.8 to the following.

Corollary 2.9. *Assume Hypothesis 1 and that for generic $b \in V$, the map $C_b \rightarrow E \times \mathbb{P}^1$ has no floating components. There is a map $V \rightarrow E_0 \times \dots \times E_0$ that sends*

$$b \mapsto \left(\varphi(\sigma_1(b)), \dots, \varphi(\sigma_n(b)) \right)$$

Let P be the image of this map. Then

$$\dim V \leq d + g - 1 + \dim P$$

If equality holds, then each component of the fiber of $V \rightarrow P$ over $(p_1, \dots, p_n) \in P$ is dense in a component of $\tilde{V}_{N'e+df, g}^{\circ, \text{tails}}(\alpha, 0)[p_1, \dots, p_k]$.

The elliptic tails arose naturally as an equality case in our dimension bound, but they do not actually occur as residual curves in the degeneration argument. To see that, we ask when elliptic tails can be smoothed away, that is, when is a map $f: C \rightarrow E \times \mathbb{P}^1$ with elliptic tails a limit of a map with smooth source. By dimensional reasons, the general point of $V_{df+N'e, g}^{\circ, \text{tails}} \setminus V_{df+N'e, g}^\circ$ cannot be smoothed into $V_{df+N'e, g}^\circ$. But more is true—as long as f is unramified, it cannot be smoothed!

Proposition 2.10. *Let $f: X = C \cup_q \tilde{E} \rightarrow E \times \mathbb{P}^1$ be a finite flat unramified map from a nodal curve such that \tilde{E} lands in a E fiber, and $\tilde{E} \rightarrow E$ is an isogeny. Then the map f cannot be deformed in a way that the node q is smoothed.*

In particular, no component of $V_{Ne+df,g}^{\circ, \text{tails}}(\alpha, 0)$ is contained in the closure of $V_{df+Ne,g}^{\circ}$.

Proof. By Theorem 5.1 of [Vis97], the deformations of the unramified map f correspond to $\text{Hom}_{\mathcal{O}_X}(\mathcal{N}_f^{\vee}, \mathcal{O}_X)$, where \mathcal{N}_f^{\vee} is the dual normal sheaf defined by:

$$(5) \quad 0 \rightarrow \mathcal{N}_f^{\vee} \rightarrow f^* \Omega_{E \times \mathbb{P}^1} \rightarrow \Omega_X \rightarrow 0$$

The key computation is the following.

Claim 2.11. *The restriction $\mathcal{N}_f^{\vee}|_{\tilde{E}}$ is equal to $\mathcal{O}_{\tilde{E}}(-q)$.*

This implies that the sheaf \mathcal{N}_f^{\vee} is locally free at q . The map to the component of the T_X^1 sheaf supported on q is just the restriction map

$$(6) \quad \text{Def}(f) = H^0(X, \mathcal{N}_f) \rightarrow H^0(X, \mathcal{N}_f|_q) = H^0(X, T_X^1|_q)$$

To establish that the node q cannot be smoothed away, we just need to show that the map in eq. (6) is zero. But the restriction map factors as

$$H^0(X, \mathcal{N}_f) \rightarrow H^0(\tilde{E}, \mathcal{N}_f|_{\tilde{E}}) \rightarrow H^0(\tilde{E}, k_q)$$

and the latter map $H^0(\tilde{E}, \mathcal{N}_f|_{\tilde{E}}) = H^0(\tilde{E}, \mathcal{O}_{\tilde{E}}(q)) \rightarrow H^0(\tilde{E}, k_q)$ is zero, since

$$H^0(\tilde{E}, \mathcal{O}_{\tilde{E}}) \xrightarrow{\sim} H^0(\tilde{E}, \mathcal{O}_{\tilde{E}}(q)) \xrightarrow{0} H^0(\tilde{E}, k_q) \xrightarrow{\sim} H^1(\tilde{E}, \mathcal{O}_{\tilde{E}})$$

This finishes the proof of Proposition 2.10, given Claim 2.11. \square

Proof of Claim 2.11. Restricting the sequence (5), we get

$$(7) \quad \text{Tor}_{\mathcal{O}_X}^1(\Omega_X, \mathcal{O}_{\tilde{E}}) \rightarrow \mathcal{N}_f^{\vee}|_{\tilde{E}} \rightarrow \pi^* \Omega_{E \times \mathbb{P}^1} \rightarrow \Omega_X|_{\tilde{E}} \rightarrow 0$$

where $\pi: \tilde{E} \rightarrow E \times \mathbb{P}^1$ is the restriction of f to \tilde{E} . We can compute the extremal terms of the sequence as follows.

Claim 2.12. *The sheaf $\text{Tor}_{\mathcal{O}_X}^1(\Omega_X, \mathcal{O}_{\tilde{E}})$ vanishes, and $\Omega_X|_{\tilde{E}} = \Omega_{\tilde{E}} \oplus k_q$.*

Proof. We only have to check this around the node q . We replace X by $\text{Spec } R$, for $R = k[x, y]/(xy)$, and \tilde{E} by the $x = 0$ branch. Then Ω_R has the free resolution

$$0 \rightarrow R\langle ydx + xdy \rangle \rightarrow R\langle dx, dy \rangle \rightarrow \Omega_R \rightarrow 0$$

Tensoring with $R/(x)$, we get the complex

$$0 \rightarrow R/(x)\langle ydx \rangle \rightarrow R/(x)\langle dx, dy \rangle$$

This is exact at $R/(x)\langle ydx \rangle$, and hence Tor^1 vanishes. The cokernel is $R/(x)\langle dy \rangle \oplus R/(x, y)\langle dx \rangle$, which implies that $\Omega_X|_{\tilde{E}} = \Omega_{\tilde{E}} \oplus k_q$. \square

We use Claim 2.12 to simplify the sequence (7), and compare it to the normal sheaf sequence of $\pi: \tilde{E} \rightarrow E \times \mathbb{P}^1$ as follows.

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathcal{N}_f^\vee|_{\tilde{E}} & \longrightarrow & \pi^* \Omega_{E \times \mathbb{P}^1} & \longrightarrow & \Omega_{\tilde{E}} \oplus k_q \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \mathcal{N}_\pi^\vee & \longrightarrow & \pi^* \Omega_{E \times \mathbb{P}^1} & \longrightarrow & \Omega_{\tilde{E}} \longrightarrow 0 \end{array}$$

The snake lemma gives us the exact sequence

$$0 \rightarrow \mathcal{N}_f^\vee|_{\tilde{E}} \rightarrow \mathcal{N}_\pi^\vee \rightarrow k_q \rightarrow 0$$

Hence, it is enough to show that $\mathcal{N}_\pi^\vee = \mathcal{O}_{\tilde{E}}$. This in turn follows from the exact sequence

$$0 \rightarrow \mathcal{N}_\pi^\vee \rightarrow \pi^* \Omega_{E \times \mathbb{P}^1} \rightarrow \Omega_{\tilde{E}} \rightarrow 0$$

since $\pi^* \Omega_{E \times \mathbb{P}^1} = \Omega_{\tilde{E}} \oplus \mathcal{O}_{\tilde{E}}$, and the map to $\Omega_{\tilde{E}}$ is the projection onto the first factor. \square

2.1. Covers of E of low gonality. In the proof of Lemma 2.5, we left two details to be checked later: that the general curve parametrized by $V_{df+Ne,g}(\alpha, 0)$ does not have tacnodes, for $d = 2$, and that it does not have triple points, for $d = 3$. Fortunately, we can deal with these cases directly through ad hoc methods, as follows.

2.1.1. The hyperelliptic locus ($d = 2$). The intersection data with E_0 is an effective divisor $D = p_1 + p_2$ on E_0 , which is fixed. This defines a degree 2 map $\varphi: E \rightarrow \mathbb{P}^1$. We will translate our problem of studying curves in $E \times \mathbb{P}^1$ to studying curves in $\mathbb{P}^1 \times \mathbb{P}^1$, using the map

$$\Phi = (\varphi, \text{id}): E \times \mathbb{P}^1 \rightarrow \mathbb{P}^1 \times \mathbb{P}^1$$

The divisor $D = p_1 + p_2$ is a fiber of φ . Let $p_0 = \Phi(p_1) = \Phi(p_2)$.

Let $C \subset E \times \mathbb{P}^1$ be a curve parametrized by $V_{df+Ne,g}(\alpha, 0)[p_1, p_2]$. Then the map Φ restricted to C is two to one. Indeed, every intersection with $E \times \{t\}$ is linearly equivalent to D , and hence a fiber of φ .

Let $\bar{C} = \Phi(C)$ be its image. It has class $(N, 1)$ in $\mathbb{P}^1 \times \mathbb{P}^1$, and is a smooth rational curve. Note that $p_0 \in \bar{C}$.

The map $\varphi: E \rightarrow \mathbb{P}^1$ has four branch points. Set

$$D = \{\text{branch point of } \varphi\} \times \mathbb{P}^1 \subset \mathbb{P}^1 \times \mathbb{P}^1$$

Away from D , the map Φ is étale. Hence, the curve C can be singular only along the preimage of D . We can relate the singularity of C with the tangency index of \bar{C} with D . We exhibit the first cases in Table 1.

We will constrain the tangency profile of \bar{C} with D . Let

$$\begin{aligned} R &= \sum_{p \in D} (\text{mult}_p(\bar{C} \cdot D) - 1) \\ &= \bar{C} \cdot D - |\bar{C} \cap D| = 4N - |\bar{C} \cap D| \end{aligned}$$

This is the total ramification of $\bar{C} \rightarrow \mathbb{P}^1$ over the four branch points of $\varphi: E \rightarrow \mathbb{P}^1$. Let E^\times (resp. $(\mathbb{P}^1)^\times$) be the complement of the ramification (resp. branch) points

Singularity of C		Tangency of \overline{C} with D	
Type	Local Equation	Multiplicity	Local Equation
Smooth	$x = y^2$	1	$x = y$
Node	$x^2 = y^2$	2	$x^2 = y$
Cusp	$x^3 = y^2$	3	$x^3 = y$
Tacnode	$x^4 = y^2$	4	$x^4 = y$

TABLE 1. A dictionary between tangency conditions of \overline{C} with D , and singularities of $C \subset E \times \mathbb{P}^1$.

of φ . Let C^\times (resp. \overline{C}^\times) be the corresponding preimage. Using Riemann–Hurwitz, we can count the ramification as follows.

$$\begin{aligned}
2g - 2 &= \text{Ram}(\varphi: C \rightarrow E) \\
&\geq \text{Ram}(\varphi: C^\times \rightarrow E^\times) \\
&= 2 \text{Ram}(\overline{C}^\times \rightarrow (\mathbb{P}^1)^\times) \\
&= 2 (\text{Ram}(\overline{C} \rightarrow \mathbb{P}^1) - R) \\
&= 2(2N - 2 - R) = 2(|\overline{C} \cap D| - 2N - 2)
\end{aligned}$$

Equality holds only if $C \rightarrow E$ is not branched over the four ramification points of φ . However, if \overline{C} had more than simple tangencies along D , then $C \rightarrow E$ would be ramified there, as we can see from the local equations in Table 1. Hence, if equalities hold, all tangencies of \overline{C} with D are simple, and the curve $C \subset E \times \mathbb{P}^1$ is nodal.

We can rewrite the inequality as

$$(8) \quad |\overline{C} \cap D| \leq 2N + g + 1$$

We now invoke Theorem 2.1 to bound the dimension in which $\overline{C} \subset \mathbb{P}^1 \times \mathbb{P}^1$ varies. Let Σ be $\text{Bl}_{p_0} \mathbb{P}^1 \times \mathbb{P}^1$, the blow up of $\mathbb{P}^1 \times \mathbb{P}^1$ at p_0 , and set $\tau = Nf_1 + f_2 - e$, where e is the exceptional divisor. We use the same D as above, and set $\Omega = \emptyset$. By eq. (8), we may set $b = 2N + g + 1$. We have

$$\begin{aligned}
\gamma &= -(K_{\text{Bl}_{p_0} \mathbb{P}^1 \times \mathbb{P}^1} + D) \cdot \overline{C} + 2N + g + 1 \\
&= -(-2f_1 + 2f_2 + e)(Nf_1 + f_2 - e) + 2N + g + 1 \\
&= -(2N - 2 + 1) + 2N + g + 1 = g + 2
\end{aligned}$$

And hence the dimension bound

$$\dim V_{Ne+2f,g}(\alpha, 0) \leq \gamma + 0 - 1 = g + 1$$

On the other hand, by Lemma 2.4, the opposite inequality holds as well. Hence, all equalities hold, including in eq. (8). As $\gamma \geq 2$, Theorem 2.1 tell us that equality must hold in (8), and hence, C is nodal, as we wanted to show.

2.1.2. *The trigonal locus* ($d = 3$). We now deal with the case of $V_{3f+Ne,g}(\alpha, 0)$. We want to show that the general point corresponds to a nodal curve, and we just have to eliminate the case of triple points.




FIGURE 2. Recovering the curve C as the intersection of $D \times \mathbb{P}^1$ and the ruled surface traced out by the lines $\varphi(t)$.

The tangency conditions with E_0 define a fixed divisor $D = p_1 + p_2 + p_3$ on E_0 . This divisor defines an embedding of $E \hookrightarrow \mathbb{P}^2$ as a smooth cubic. Each intersection $C \cap (E \times \{t\})$ maps to three collinear points in $E \subset \mathbb{P}^2$. We think of this as a map $\varphi: \mathbb{P}^1 \rightarrow \widehat{\mathbb{P}^2}$ sending t to the line joining the three collinear images of $C \cap (E \times \{t\})$. Note that $\varphi(0)$ goes to the fixed line ℓ_0 connecting the images of p_1, p_2 and p_3 .

Conversely, given the map φ , we can reconstruct C as the intersection of $E \times \mathbb{P}^1 \subset \mathbb{P}^2 \times \mathbb{P}^1$ with the ruled surface in $\mathbb{P}^2 \times \mathbb{P}^1$ traced out by the lines $\varphi(t)$ as we vary $t \in \mathbb{P}^1$. See fig. 2.

We can relate the singularities of C with the tangency of the image of φ along $\widehat{E} \subset \widehat{\mathbb{P}^2}$, the dual to $E \subset \mathbb{P}^2$. The dual \widehat{E} is a degree 6 curve with 9 cusps (one for each flex of E).

If C is singular at $(p, t) \in E \times \mathbb{P}^1$, then $\varphi(t)$ lies on \widehat{E} —specifically, at the point corresponding to the tangent to E at p . Conversely, if we denote by $\tilde{C} \rightarrow C$ the normalization map, we have:

- If $\varphi(\mathbb{P}^1)$ is transverse to \widehat{E} at a point $\varphi(t)$, then $\tilde{C} \rightarrow \mathbb{P}^1$ is simply branched over t , and C is smooth at the corresponding point of $E \times \mathbb{P}^1$.
- If $\varphi(\mathbb{P}^1)$ is simply tangent to \widehat{E} at a point $\varphi(t)$, then $C \subset E \times \mathbb{P}^1$ is nodal at the corresponding point, and $\tilde{C} \rightarrow \mathbb{P}^1$ is not branched over t .
- If $\varphi(\mathbb{P}^1)$ does not contain any of the 9 cusps of \widehat{E} , then C has no triple points.

Hence, it is enough to show that the image of $\varphi(\mathbb{P}^1)$ has at most simple tangencies with \widehat{E} , and does not pass through any of the cusps. We will establish that by invoking once more Theorem 2.1, now to bound the dimension in which the image of φ can vary.

By Riemann–Hurwitz, $\widetilde{C} \rightarrow \mathbb{P}^1$ has $2g + 4$ branch points. Let T be the number of transverse intersections of $\varphi(\mathbb{P}^1)$ and \widehat{E} . Hence $T \leq 2g + 4$. Moreover,

$$\begin{aligned} 6N &= \varphi(\mathbb{P}^1) \cdot \widehat{E} \\ &= \sum_p \text{mult}_p(\varphi(\mathbb{P}^1) \cdot \widehat{E}) \\ &\geq T + 2(|\varphi(\mathbb{P}^1) \cap \widehat{E}| - T) \\ &= 2|\varphi(\mathbb{P}^1) \cap \widehat{E}| - T \\ &\geq 2|\varphi(\mathbb{P}^1) \cap \widehat{E}| - (2g + 4) \end{aligned}$$

With equality holding only if at most simple tangencies occur. Rearranging, we get

$$(9) \quad 3N + g + 2 \geq |\varphi(\mathbb{P}^1) \cap \widehat{E}|$$

We set $\Sigma = \text{Bl}_{\ell_0} \widehat{\mathbb{P}^2}$. The class τ of the proper transform of $\varphi(\mathbb{P}^1)$ is $Nh - e$, where e is the class of the exceptional divisor. Using the notation of Theorem 2.1, and setting $b = 3N + g + 2$, we have

$$\begin{aligned} \gamma &= -(K_\Sigma + \widehat{E}) \cdot \tau + 3N + g + 2 \\ &= -(3h + e)(Nh - e) + 3N + g + 2 \\ &= -(3N + 1) + 3N + g + 2 = g + 1 \end{aligned}$$

So the image of $\varphi(\mathbb{P}^1)$ can vary in at most $0 - 1 + g + 1 = g$ dimensions. Adding $\dim(\text{Aut}(\mathbb{P}^1, 0)) = 2$ to remember the map $\varphi: \mathbb{P}^1 \rightarrow \widehat{\mathbb{P}^1}$ (we know that 0 maps to ℓ_0), we get that the dimension of $V_{3f+Ne,g}(\alpha, 0)$ is at most $g + 2$. However, by Lemma 2.4, the opposite inequality holds as well. Hence, equality in (9) holds, and the original curve C is nodal.

3. HYPERPLANE SECTIONS OF $\widetilde{V}_{Ne+df,g}^\circ(a, b)$

In this section, we will prove Theorems 1.18 and 1.20 stated in the introduction, and Theorem 3.13 which is a generalization of them. Let us briefly remind ourselves of the setup. We fix a fiber $E_0 \subset E \times \mathbb{P}^1$, and let $p \in E_0$ be a general point. We want to characterize the irreducible components of $\widetilde{V}_{Ne+df,g}^\circ(a, b) \cap H_p$. Let W be one of these components, and $Y_0 \subset E \times \mathbb{P}^1$ be the curve corresponding to a general point in W .

If the curve Y_0 does not contain E_0 , it is easy to describe what W must be—we will do that in Section 3.4. Let us assume, then, that the curve E_0 is contained in Y_0 . Now it is much harder to analyze the residual curve. Our main tool to get our hands on its geometry is the following construction.

3.1. Nodal reduction. Pick a general point $Y_0 \subset E \times \mathbb{P}^1$ in W . Choose a general arc Δ in $\tilde{V}_{Ne+df,g}^\circ(a, b)$ approaching the point $(Y_0) \in W$. We may assume that Δ is smooth by taking its normalization.

We get a flat family of curves mapping to $E \times \mathbb{P}^1$,

$$\begin{array}{ccc} \mathcal{Y} & \xrightarrow{\varphi} & E \times \mathbb{P}^1 \\ \downarrow \pi & & \\ \Delta & & \end{array}$$

where the generic fiber is a general point in $\tilde{V}_{Ne+df,g}^\circ(a, b)$, and the central fiber is a generic point in W . Let \mathcal{Y}^\times be the restriction of \mathcal{Y} over the punctured base Δ^\times .

The preimage of $E_0 \subset E \times \mathbb{P}^1$ under φ on the punctured family \mathcal{Y}^\times will consist of some multisections. After a base change, we may assume that the multisections are actually sections. Being in $\tilde{V}_{Ne+df,g}^\circ(a, b)$, there will be a of those which get contracted under the map φ , and b that are allowed to vary—that is, are not necessarily contracted by φ . To avoid overburden our notation even more, we will refer to these as A -sections and B -sections.

At this point there is not much we can say about the central fiber. However, if we perform nodal reduction on $\mathcal{Y} \rightarrow \Delta$, we can replace it by a more tractable object. Let us denote by $\mathcal{C} \rightarrow B$ the family obtained by nodal reduction, and $f: B \rightarrow \Delta$ the corresponding base change map. Here is what we can assume about it.

Hypothesis 2. The family $\mathcal{C} \rightarrow B$ satisfies the following properties.

- The total space \mathcal{C} is smooth.
- The fibers C_b are equal to the corresponding fiber $Y_{f(b)}$ of the original family $\mathcal{Y} \rightarrow \Delta$, as long as $b \neq 0$.
- The (scheme-theoretic) central fiber C_0 is nodal (and therefore reduced).
- The maps $\varphi: C_b \rightarrow E \times \mathbb{P}^1$ are semistable.
- The preimage of E_0 is a union of sections of $\mathcal{C} \rightarrow B$, plus components in the central fiber. These sections are disjoint on the general fiber, but are allowed to meet along the central fiber. We will still call them A -sections and B -sections, as above.
- The family $\mathcal{C} \rightarrow B$ is the minimal one respecting the conditions above.

See fig. 3 for a diagram of the family $\mathcal{C} \rightarrow B$ away from the central fiber. Our goal is to determine how the central fiber looks like, and where the A and B -sections meet it.

The idea of the proof of Theorem 1.18 is simple, but combinatorially challenging. There are two opposing forces controlling the central fiber. On one hand, it varies in a large dimension, and hence its singularities should be mild, and few extra conditions should be imposed. That is, the dimension count forces the central fiber to be “simple”. On the other hand, from the parametric point of view, the flatness of the family \mathcal{C} says the central fiber has to have arithmetic genus g , and hence it cannot be *too* simple. We will determine the equilibrium positions, the conditions under which the two forces precisely cancel out. And it will turn out that these conditions completely describe the components of the hyperplane section.




FIGURE 3. A diagram of the family $\mathcal{C} \rightarrow B$ away from the central fiber.

To make the structure of the proof transparent, we will break up the argument into a chain of small claims. Each one will be an inequality expressing the tension between the two points of view. As we compose this long chain of inequalities, we will conclude that all equalities must hold! We then come back and use the equality conditions to pin down the possible components of the hyperplane section.

Again, we do not prove that all of these components do appear in a hyperplane section, or even if they do appear, we do not say with what multiplicity. Here we only compile a list of *suspects*: any component W must be in this list.

The difficulty is that our starting point, the central fiber C_0 , is an arbitrary nodal curve. Just notionally this is already very cumbersome. We have to design our notation to express all the arguments in this generality.

We divide the components of C_0 in three groups, forming nodal curves each:

- We call \tilde{E} the curve which is the union of irreducible components of C_0 that dominate E_0 under the map $\varphi: C_0 \rightarrow E \times \mathbb{P}^1$.
- Let Z be the union of the components that get contracted to a point in E_0 by the map φ_0 .
- Let X be the union of the remaining components.

Let T be the number of connected components of Z meeting both X and \tilde{E} , plus the number of nodes between X and \tilde{E} . For an example, see fig. 4.

3.2. The genus bound. In this section we will leverage the flatness of the family $\mathcal{C} \rightarrow B$ to bound the arithmetic genus of X in terms of g and T .

Start by enumerating the irreducible components of Z by Z_1, \dots, Z_m . To help us keep track of the combinatorial data of C_0 , we introduce an auxiliary graph G . We obtain G by starting with the *dual graph* of the nodal curve C_0 , and then




FIGURE 4. An example of an a priori possible central fiber C_0 , with its components grouped as \tilde{E} , Z and X . In this example, $T = 4$. The central fibers that do arise will turn out to be much simpler than this, as we will see.

collapsing the subgraph corresponding to components of X (resp. \tilde{E}) into a single vertex. Hence, it has $m + 2$ vertices—one for each Z_i , and one for \tilde{E} and one for X , even though \tilde{E} and X may well be reducible. The edges of G correspond to the nodes of C_0 which are not *internal* to X or \tilde{E} —that is, we don't consider nodes which connect two components of X , or two components of \tilde{E} . In particular, there are no loops in G at the vertices corresponding to \tilde{E} and X . We will refer to the nodes corresponding to edges in G as *external* nodes. Note that the number of edges in G is equal to the number of external nodes. For an example, see fig. 5.

Claim 3.1. *Let d_X (resp. $d_{\tilde{E}}$) be the degree of X (resp. \tilde{E}) in the graph G . Then, $T \leq d_X$ and $T \leq d_{\tilde{E}}$. If both equalities hold, then every connected component of Z meets X and \tilde{E} exactly once each.*

Proof. This follows directly from the definition of T . The inequality is strict in two situations: if there are connected components of Z meeting only one of X and \tilde{E} , or if there are connected components of Z meeting X or \tilde{E} in more than one node. \square

Claim 3.2. *Let d_{Z_i} be the number of nodes on Z_i in C_0 , that is, the degree of Z_i in the graph G . Then*

$$0 \leq g(Z_i) - 1 + \frac{d_{Z_i}}{2}$$

with equality only if Z_i is rational with exactly two nodes on it.

Proof. This comes from the minimality of the family. If this quantity was negative, then Z_i would be a rational curve meeting the rest of the central fiber at only one




FIGURE 5. The dual graph of the example C_0 in fig. 4, and the respective graph G obtained by contracting the subgraphs corresponding to X and \tilde{E} .

point. Then, we compute the intersection pairing on the surface \mathcal{C} :

$$0 = Z_i C_0 = Z_i^2 + 1$$

from which we conclude that Z_i is a (-1) -curve. Moreover, being a component in Z , the map φ contracts Z_i . Hence, we could contract it in \mathcal{C} and maintain the conditions of Hypothesis 2, contradicting the minimality of \mathcal{C} .

If equality holds, then $g(Z_i) = 0$ and $d_{Z_i} = 2$, as we want, or $g(Z_i) = 1$ and $d_{Z_i} = 0$. The latter case does not occur, because Z_i would be a disconnected component of C_0 which gets contracted under the map to $E \times \mathbb{P}^1$, and hence cannot be in the limit of maps $C_t \rightarrow E \times \mathbb{P}^1$ which are birational onto their image. \square

Next, we compute the arithmetic genus of the central fiber. By summing the degrees of all vertices of G , we get twice the number of edges. That is,

$$d_X + d_{\tilde{E}} + \sum d_{Z_i} = 2 \times |\text{edges in } G| = 2 \times |\text{external nodes in } C_0|$$

Hence, we get

$$\begin{aligned} p_a(C_0) - 1 &= p_a(X) - 1 + p_a(\tilde{E}) - 1 + \sum_{i=1}^m (g(Z_i) - 1) + |\text{external nodes in } C_0| \\ &= p_a(X) - 1 + p_a(\tilde{E}) - 1 + \frac{d_X + d_{\tilde{E}}}{2} + \sum_{i=1}^m \left(g(Z_i) - 1 + \frac{d_{Z_i}}{2} \right) \end{aligned}$$

Moreover, by flatness,

$$g - 1 = p_a(C_t) - 1 = p_a(C_0) - 1$$

Summing up, we get the following.




FIGURE 6. A picture of C_0 in a case where equality in the genus bound (Proposition 3.5) holds.

Claim 3.3. *We have,*

$$g - 1 = p_a(X) - 1 + p_a(\tilde{E}) - 1 + \frac{d_{\tilde{E}} + d_X}{2} + \sum_{i=1}^m \left(g(Z_i) - 1 + \frac{d_{Z_i}}{2} \right)$$

The last ingredient is the following.

Claim 3.4. *We have*

$$0 \leq p_a(\tilde{E}) - 1$$

with equality only if \tilde{E} is a disjoint union of smooth genus 1 curves, and $\tilde{E} \rightarrow E_0$ is unramified.

Proof. This follows from the Riemann–Hurwitz formula applied to $\tilde{E} \rightarrow E_0$. Note that \tilde{E} may be disconnected. \square

We are ready for the following key bound.

Proposition 3.5. *[The Genus Bound] We have*

$$p_a(X) + T \leq g$$

Equality holds only if \tilde{E} is smooth, $\tilde{E} \rightarrow E$ is unramified, and Z is a union of chains of rational curves connecting X and \tilde{E} .

The equality does not hold in our previous example fig. 4, but in fig. 6 we can see an example where it does hold.

Proof. Just add up the inequalities in Claims 3.1 to 3.4. If equality holds, then it does as well in every claim, and hence all equalities conditions are true.

From Claim 3.4 we get that \tilde{E} is smooth and $\tilde{E} \rightarrow E$ is unramified. The equality conditions of Claim 3.2 imply that Z is a disjoint union of chains of rational

curves. Claim 3.1 implies that each of these chains connects a component of \tilde{E} to a component of X . \square

Note that, so far, even if equality occurs in Proposition 3.5, we can't say much about the residual curve X . Also note that we haven't used the sections our family came equipped with either. These will play a role in the following.

3.3. The dimension bound. We now want to constraint the geometry of X based on the fact it varies in a family of dimension $\dim W = \dim H_p \cap \tilde{V}_{Ne+df,g}^\circ(a, b) = d + g + b - 3$.

The construction in Section 3.1 replaces a general point $\{Y_0 \subset E \times \mathbb{P}^1\}$ of the component W with a map $C_0 \rightarrow E \times \mathbb{P}^1$. To leverage the dimension of W , we need to extend this *pointwise* construction to a *local* one. The following technical lemma does so.

Lemma 3.6. *For each general point $p \in W$, we can find a map $\eta: \mathcal{U} \rightarrow \mathcal{U}_0 \subset W$ which is finite onto a open subscheme which contains p , a flat family $\mathcal{C}^{\mathcal{U}} \rightarrow \mathcal{U}$ and a map $\varphi: \mathcal{C}^{\mathcal{U}} \rightarrow E \times \mathbb{P}^1$ such that for each $u \in \mathcal{U}$, the map $C_u^{\mathcal{U}} \rightarrow E \times \mathbb{P}^1$ is equal to the result of the pointwise construction of Section 3.1 on $\eta(u) \in W$.*

Proof. The idea is the following: the only choice we had in the construction in Section 3.1 was of the arc $\Delta \rightarrow V$. However, as we are performing a codimension one degeneration, the limit will not depend on the direction of approach. Here is a standard way to formalize this.

The variety $\tilde{V}_{Ne+df,g}^\circ(a, b)$ admits a rational map to the Kontsevich space of stable maps $\overline{M}_{g,a+b}(E \times \mathbb{P}^1, df + Ne)$. Since the target is projective and the source is normal, we can extend the rational map in codimension one. Hence, for a general point in the divisor $H_p \cap \tilde{V}_{Ne+df,g}^\circ(a, b)$, we can find a neighborhood of p mapping to the coarse space $\overline{M}_{g,a+b}(E \times \mathbb{P}^1, df + Ne)$. Possibly after a base change, we find some neighborhood $\hat{\mathcal{U}}$ which maps to the stack $\overline{\mathcal{M}}_{g,a+b}(E \times \mathbb{P}^1, df + Ne)$. Now pulling back the universal family, we get a family $\mathcal{C}^{\hat{\mathcal{U}}} \rightarrow \hat{\mathcal{U}}$. Let $\mathcal{U} \subset \hat{\mathcal{U}}$ be the preimage of W . Our universal family $\mathcal{C}^{\mathcal{U}} \rightarrow \mathcal{U}$ is the restriction of $\mathcal{C}^{\hat{\mathcal{U}}} \rightarrow \hat{\mathcal{U}}$. \square

The total space $\mathcal{C}^{\mathcal{U}}$ is reducible: there is at least a component containing the generic fiber's \tilde{E} , and a component containing the generic fiber's X . We want to pick out the component $\mathcal{X} \subset \mathcal{C}^{\mathcal{U}}$ corresponding to the X 's. A way to directly construct it is to set

$$\mathcal{X} = \overline{\mathcal{C} - \varphi^{-1}(E_0)}, \text{ where } \varphi: \mathcal{C}^{\mathcal{U}} \rightarrow E \times \mathbb{P}^1$$

Let us denote the restriction of $\varphi: \mathcal{C}^{\mathcal{U}} \rightarrow E \times \mathbb{P}^1$ to \mathcal{X} by φ again. We will bound arithmetic genus $p_a(X_u)$ in terms of the dimension of \mathcal{U} and the tangency conditions the image of the X_u satisfies along E_0 , using Corollary 2.9.

To get our hands on the tangency conditions, we need to introduce some notation. Let k be the cardinality of $\varphi^{-1}(E_0) \cap X_u$ for u generic. After a base change, we may assume that $\varphi^{-1}(E_0)$ consists of k distinct sections. Let $\mathcal{U} \rightarrow E_0^k$ the map sending u to the images under φ of the k sections. Let $P \subset E_0^k$ be the image of this map.

Each one of the k sections appears with some multiplicity, say α_i . We get a corresponding map

$$\begin{aligned} E_0^k &\rightarrow \text{Div}^d E_0 \\ (p_1, \dots, p_k) &\mapsto \sum \alpha_i p_i \in \text{Div}^d E_0 \end{aligned}$$

For example, the composite map $U \rightarrow E_0^k \rightarrow \text{Div}^d(E_0)$ sends u to the divisor $\varphi(X_u) \cdot E_0$ on E_0 .

Note that the linear class of $\varphi(X_u) \subset E \times \mathbb{P}^1$ is equal to the class of C_0 minus some number of E_0 's that got split off. That is, $[\varphi(X_u)] = \mathcal{L} - m e$, and the intersection $\varphi(X_u) \cdot E_0 = \mathcal{L}|_{E_0}$ is a fixed linear class. Hence, the following diagram commutes.

$$\begin{array}{ccccccc} U & \longrightarrow & P \subset E_0^k & \longrightarrow & \text{Div}^d E_0 & \longrightarrow & \text{Pic}^d E_0 \\ & \searrow \text{dashed} & & & \uparrow & & \uparrow \\ & & & & |\mathcal{L}|_{E_0}| & \longrightarrow & \text{Spec } \mathbb{C} = \{\mathcal{L}|_{E_0}\} \end{array}$$

Denote by X the generic fiber X_u .

Claim 3.7. *We have*

$$\dim W \leq d + p_a(X) - 1 + \dim P$$

If equality holds, then, after a base change, W is fibered over P , and the fiber over $(p_1, \dots, p_k) \in P \subset E_0^k$ is dense in a component of

$$\widetilde{V}_{df+N'e, p_a(X)}^{\circ, \text{tails}}(\alpha, 0)[p_1, \dots, p_k]$$

where α_i is the multiplicity in which the section over the point p_i appears, and $df + N'e$ is the homology class of $\varphi(X)$.

Proof. This follows from Corollary 2.9 applied to $\mathcal{X} \rightarrow U$, and the fact that U and W have the same dimension. \square

To bound the dimension of P , let us go back to our one-dimensional degeneration $\mathcal{C} \rightarrow B$ as in Section 3.1. Let ε_b be the number of B -sections which meet the central fiber on \widetilde{E} or on a connected component of Z that does not meet X . That is, out of the b B -sections that the family $\mathcal{C} \rightarrow B$ has, $b - \varepsilon_b$ of them land in X , or in a component of Z that contracts to a point in X . Either way, these $b - \varepsilon_b$ points are in the intersection of $\varphi(X)$ with E_0 .

Define

$$\varepsilon = \begin{cases} 1, & \text{if } \varepsilon_b > 0 \\ 0, & \text{else.} \end{cases}$$

For example, we have

$$(10) \quad \varepsilon \leq \varepsilon_b$$

with equality only if ε_b is equal to zero or one.

Now we are ready to state the bound on $\dim P$.

Claim 3.8. *We have*

$$\dim P \leq T + b - \varepsilon_b - 1 - (1 - \varepsilon)$$

If equality holds, none of the A or B-sections land in components of Z meeting X and \tilde{E} , and one of the two following scenarios holds.

- (1) *If $\varepsilon = 1$, then P is dense in the preimage under $E_0^k \rightarrow \text{Div}^d E_0$ of the sum of:*
 - *a full linear series of degree $T + b - \varepsilon_b$,*
 - *and $\alpha' = d - (T + b - \varepsilon_b)$ base points.*
- (2) *Else, both ε_b and ε are zero, and P is dense in the preimage under $E_0^k \rightarrow \text{Div}^d E_0$ of the sum of:*
 - *a full linear series of degree b on E (corresponding to the limit of the B-sections),*
 - *a full linear series of degree T on E ,*
 - *and $\alpha' = d - b - T$ fixed base points.*

Proof. The curve $\varphi(X)$ meets E_0 in three types of points:

- The image of A-sections. These add no moduli to P .
- The images of B-sections. There are $b - \varepsilon_b$ of those.
- The images of connected components of Z meeting both X and \tilde{E} , or nodes connecting \tilde{E} and X in C_0 . There are exactly T of those.

Hence, the dimension of P is bounded by

$$\dim P \leq T + b - \varepsilon_b$$

However, if equality were to hold, then P would be dense in $E^{T+b-\varepsilon_b} \subset E^k$ for some subset of the indices. But $P \rightarrow E^k \rightarrow \text{Div}^d E$ has to land inside the full linear series $|\mathcal{L}|_{E_0} \subset \text{Div}^d E$, while the image of $E^{T+b-\varepsilon_b} \subset E^k \rightarrow \text{Div}^d E$ does not have a fixed linear class. This is a contradiction.

Hence, we get the slightly stronger inequality

$$\dim P \leq T + b - \varepsilon_b - 1$$

for which equality happens exactly if P is the preimage of some full linear series of degree $T + b - \varepsilon_b$, plus some fixed base points (which correspond to A-sections landing in X). That is, only in scenario (1).

Something interesting happens when $\varepsilon_b = 0$. In this case, all the B-sections land in X , and their sum is a fixed divisor up to linear equivalence. Indeed, it is exactly the class of $\varphi(C_t) = \mathcal{L}$, restricted to E_0 , minus the fixed images of the A-sections p_1, \dots, p_a . Hence, we get one moduli less than expected. That is, we may subtract $(1 - \varepsilon)$ from the count above, to get to

$$\dim P \leq T + b - \varepsilon_b - 1 - (1 - \varepsilon)$$

as we wanted to prove. Now equality happens in both scenarios (1) and (2). \square

Putting all of this together, we get the following.

Proposition 3.9 (The dimension bound). *We have*

$$g \leq p_a(X) + T$$

If equality holds, then there is a partition τ of length T , and an integer $a' \leq a$, such that X is a general point of

- $\tilde{V}_{df+N'e, p_a(X)}^{\circ, \text{tails}}(a', b, \tau)$, if $\varepsilon = 0$, or
- $\tilde{V}_{df+N'e, p_a(X)}^{\circ, \text{tails}}(a', \tau + (b - 1))$, if $\varepsilon = 1$.

Proof. By Theorem 2.6, we have

$$\dim W = \dim \tilde{V}_{Ne+df, g}^{\circ}(a, b) - 1 = d + g + b - 3$$

Now combine Equation (10) and Claims 3.7 and 3.8, and we get the desired inequality.

If equality holds, we let $\tau = (\tau_1, \tau_2, \dots)$, where τ_i is the multiplicity in which each of the corresponding T points appear in the divisor $\varphi^*(E_0)|_X$, and a' be the number of A -sections landing in X . Considering all the equality conditions together, we get the description above. \square

3.4. Proof of Theorems 1.18 and 1.20. Let us start with Theorem 1.18. Let W be an irreducible component of $\tilde{V}_{Ne+df, g}^{\circ}(a, b) \cap H_p$. If the generic point of W is a curve that does not contain E_0 , then W is contained in

$$\tilde{V}_{Ne+df, g}^{\circ}(a + 1, b - 1)[p_1, \dots, p_a, p]$$

However, by Theorem 2.6, this has the same dimension as W . Hence, W must be one of its irreducible components, as we wanted to show.

Let us assume now that the general point of W is a curve which contains E_0 with multiplicity $m > 0$. Set $N' = N - m$, so that the homology class of the residual curve, $\varphi(X)$, is $N'e + df$.

Then we apply the results of Sections 3.1 to 3.3. On one hand, Proposition 3.5 (the genus bound) says

$$p_a(X) + T \leq g$$

while Proposition 3.9 (the dimension bound) says

$$g \leq p_a(X) + T$$

Hence, both equalities have to hold! We can now use the equality conditions to describe the central fiber C_0 and the way the A and B -sections meet it.

By Proposition 3.5, the map $\tilde{E} \rightarrow E$ is unramified, and Z is just a disjoint union of T' rational chains connecting X and \tilde{E} , plus $(T - T')$ nodes between X and \tilde{E} . Let τ be a partition keeping track of the multiplicities with which each of the T rational chains/nodes connecting X and \tilde{E} appears in the pullback $\varphi^*(E_0)|_X$.

By Proposition 3.9, none of the A and B -sections hit the central fiber at Z . Let a' be the number of A -sections that land in X . In particular, $a' \leq a$.

There are two possible scenarios for X , according to if there are any B -sections approaching \tilde{E} in the limit.

If there are B -sections approaching \tilde{E} , then there is exactly one such. In this case, the residual curve X is a general point of a component of the Severi variety $\tilde{V}_{N'e+df, p_a(X)}^{\circ, \text{tails}}(a', \tau + (b - 1))$.

On the other hand, if there are no B -sections approaching X , then X is a general point of a component of $\tilde{V}_{N'e+df, p_a(X)}^{\circ, \text{tails}}(a', b, \tau)$. See figs. 7 and 8 for diagrams for the family $\mathcal{C} \rightarrow B$ and its sections.




FIGURE 7. An example of a central fiber, and of the first of the two possible scenarios of how the sections meet the central fiber—never on the Z components, and none of the B -sections meets \tilde{E} .

There is one last issue to be taken care of: we have to eliminate the elliptic tails. This follows directly from Proposition 2.10. Note that this also rules out $\tau = (1)$ —in this case, \tilde{E} itself would be an elliptic tail, which can't be smoothed. This concludes the proof of Theorem 1.18.

Remark 3.10. As opposed to the analogue arguments applied on rational surfaces [CH98, Vak00], our directrix E_0 can be split off generically with multiplicity greater than one. Compare with [SS13] where the “directrix” is still split off with multiplicity one, but the residual curve can be non-reduced.

To prove Theorem 1.20, we note that the nodal reduction construction in Section 3.1 already resolved the map $\tilde{V}_{N_{e+df,g}}^\circ(a, b) \dashrightarrow \overline{\mathcal{M}}_g(E, N)$ at the general point of W . By composing the map $\varphi: C_0 \rightarrow E \times \mathbb{P}^1$ with the projection $E \times \mathbb{P}^1 \rightarrow E$, we get a map that is almost an element of $\overline{\mathcal{M}}_g(E, N)$ —we just have to contract the rational chains Z . The description in Theorem 1.20 follows.

3.5. The general case. Let us tackle the problem of computing the degree of the Severi variety $V_{N_{e+df,g}}^\circ$. The strategy is to intersect it with multiple hyperplanes of the form H_p , for general $p \in E_0$, until we get down to some number of points, which we could count. Let V be a component that arises in the process of intersecting




FIGURE 8. The same example of a central fiber as in fig. 7, but with the second possibility for how the sections meet the central fiber—again never on the Z components, and with exactly one of the B -sections meeting \tilde{E} .

$\tilde{V}_{N'e+df,g}^\circ(a, b)$ with hyperplanes. To succeed, we have to answer the following two questions:

- what are the components of the intersection $V \cap H_p$?
- with what multiplicity does each component appear?

We will deal with only the first question. Theorem 1.18 describes the components of the intersections until the first time the directrix E_0 is split off generically. To continue, we would need an analogue of Theorem 1.18 for varieties like $\tilde{V}_{N'e+df,p_a(X)}^\circ(a', b, \tau)$.

Fortunately, the argument above readily adapts to this more general situation, and all the components that ever arise in this process can be described in similar tangency terms. We will follow an analogue of Caporaso–Harris notation for these new components. The key difference in our case is that, since E has non trivial Picard group, we are also forced to remember the classes of some of the moving points, as we already had in the case of $\tilde{V}_{N'e+df,p_a(X)}^\circ(a', b, \tau)$.

Definition 3.11. Fix line bundles $\mathcal{L}, L_j \in \text{Pic } E_0$, points $p_1, \dots, p_k \in E_0$, and tangency profiles α and $\beta^1, \dots, \beta^\ell$, such that $m(\beta^j) = \deg L_j$ and

$$\mathcal{L} = \sum_{i=1}^k \alpha_i p_i + \sum_{j=1}^{\ell} L_j$$

where $\alpha = (\alpha_1 \dots \alpha_k)$. We define the generalized Severi variety

$$\tilde{V}_{Ne+df,g}^\circ(\alpha, \beta^1, \beta^2, \dots, \beta^\ell)[p_1, \dots, p_k, L_1, \dots, L_\ell]$$

as the normalization of the closure of the locus of reduced (possibly reducible) curves in $E \times \mathbb{P}^1$, of class $\pi_1^* \mathcal{L} \otimes \pi_2^*(\mathcal{O}_{\mathbb{P}^1}(d))$, geometric genus g , containing no fibers of $E \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$, and whose intersection with E_0 consists of $|\alpha| + \sum_j |\beta^j|$ points as follows:

- The fixed points p_1, \dots, p_k , each appearing with the corresponding multiplicity $\alpha_1, \dots, \alpha_k$ (note that $|\alpha| = k$),
- and for each j , $|\beta^j|$ points with multiplicities described by β^j , and whose sum is linearly equivalent to L_j .

We will omit the brackets when possible, and just write $\tilde{V}_{Ne+df,g}^\circ(\alpha, \beta^1, \dots, \beta^\ell)$.

Note that by an argument similar to Theorem 2.6, we have

$$(11) \quad \dim \tilde{V}_{Ne+df,g}^\circ(\alpha, \beta^1, \dots, \beta^\ell) = d + g + \sum_j |\beta^j| - 1 - \ell$$

and its generic point corresponds to a nodal curve.

Remark 3.12. If for some j we have $\beta^j = (b)$, then the generalized Severi variety is a disjoint union of b^2 different generalized Severi varieties, where we remove β^j and add an extra $\alpha_{k+1} = b$, with the corresponding point p_{k+1} being a b -th root of L_j . Hence, we may assume that $|\beta^j| \geq 2$ for all j .

We can now describe the components that will ever arise in the process of intersecting with the hyperplanes H_p 's.

Theorem 3.13. *Let p be a generic point on E_0 , and W a component of the intersection $\tilde{V}_{Ne+df,g}^\circ(\alpha, \beta^1, \dots, \beta^\ell) \cap H_p$. Assume that $|\beta_j| \geq 2$ for all j . Then one of the following holds.*

- Either W is a component of

$$\tilde{V}_{Ne+df,g}^\circ(\alpha + (n), \beta^1, \dots, \beta^j - (n), \dots, \beta^\ell)[p_1, \dots, p_k, p, L_1, \dots, L_j - p, \dots, L_\ell]$$

for some j . That is, one of the moving points became a fixed at p .

- Or the elliptic fiber E_0 is generically split off in the limit with multiplicity m . As a parameter space for the residual curve, W is a component of

$$\tilde{V}_{(N-m)e+df,g'}^\circ(\alpha', \beta^{j_1}, \dots, \beta^{j_t}, \beta')[p_{i_1}, \dots, p_{i_{|\alpha'|}}, L_{j_1}, \dots, L_{j_t}, L']$$

where

- (1) $J = \{j_1, \dots, j_t\} \subset [\ell]$ is a subset
- (2) for $j \notin J$, let $\widehat{\beta^j} \subset \beta^j$ be a tangency profile with one less term.
- (3) τ is a tangency profile such that $|\tau| \geq 2$ and

$$g = g' + |\tau|$$

- (4) β is a tangency profile such that

$$\beta' = \tau + \sum_{j \notin J} \widehat{\beta^j}$$

(5) and $\alpha' \subset \alpha$ is a subprofile such that

$$d = m(\alpha') + m(\beta') + \sum_{j \in J} m(\beta^j)$$

That is, for each group β^j of B-sections, at most one of them may specialize to the curve dominating E_0 . All the others stay in the residual curve. If the group of β^j remains intact, then its class is preserved. Otherwise, all the β^j that lost an element bundle up together, plus a partition τ that remembers how the residual curve meets the cover of E_0 in the stable limit.

Proof. The argument will be very similar to the one for the simpler case of Theorem 1.18. As a matter of fact, the proof goes through mostly word by word—the only difference is that we have to do more bookkeeping to keep track of all the tangency profiles α, β^j . Let us just highlight the differences, and leave the details to the reader.

Sections 3.1 and 3.2 go through as they are. In particular, we will use the same definition of \tilde{E}, Z, X and T . We obtain Proposition 3.5 (the genus bound):

$$p_a(X) + T \leq g$$

In Section 3.3, the key adjustment is our definition of ε_b and ε . We let ε_j be the number of sections in the β^j group which meet the central fiber on \tilde{E} or on a connected component of Z that does not meet X . We define

$$\varepsilon = |\{j \text{ such that } \varepsilon_j \neq 0\}|$$

From the definition of ε , we get the following.

Claim 3.14. *We have*

$$\varepsilon \leq \sum_{j=0}^{\ell} \varepsilon_j$$

with equality only if, for each j , ε_j is equal to zero or one.

And the corresponding Claim 3.8 is the following.

Claim 3.15. *We have*

$$\dim P \leq T + \sum_j (|\beta^j| - \varepsilon_j) - (\ell - \varepsilon + 1)$$

We get the dimension bound

$$g \leq p_a(X) + T$$

by coupling Claims 3.7, 3.14 and 3.15, and the following consequence of eq. (11):

$$\dim W = \dim V_{df+N_e, g}(\alpha, \beta) - 1 = d + g + \left(\sum |\beta^j|\right) - 1 - \ell - 1$$

Again, the dimension bound balances exactly the genus bound, and equalities hold throughout. Analyzing the equality conditions, and eliminating elliptic tails by Proposition 2.10, we arrive at the list of Theorem 3.13. \square

4. MONODROMY GROUPS OF COVERS

In this section we will study the monodromy groups of simply branched covers.

Let $f: C \rightarrow D$ be a degree d cover of smooth connected curves. Let $D^\times \subset D$ be the complement of the set of branch points of f , and C^\times its preimage in C . We will deal with fundamental groups, so let us fix base points $\star \in D^\times$ and $\star \in C^\times$ that map to each other under f .

The map $f: C^\times \rightarrow D^\times$ is a covering of topological spaces. The fundamental group $\pi_1(D^\times, \star)$ acts on the fiber $f^{-1}(\star)$. This action defines a *monodromy map* $\varphi: \pi_1(D^\times, \star) \rightarrow S_d$. We call the image of φ the monodromy group of the cover f . Note that it is only defined up to conjugation in S_d . If φ is surjective, we say that f has *full monodromy*.

A driving question in the field is *what are the possible monodromy groups* G . There are many interesting (and some unsolved) problems along this line, see for example, [Gur03, AP05, FM01, GMS03]. However, when the covering map f is simply branched, one can say much more.

Proposition 4.1 (Bernstein–Edmonds). *Let $f: C \rightarrow D$ be a simply branched cover. Then f has full monodromy if, and only if, f is primitive (that is, the pushforward map in fundamental groups $\pi_1(C, \star) \rightarrow \pi_1(D, \star)$ is surjective).*

This is Proposition 2.5 of [BE84]. One can find a proof of this in [Kan05b] as well. For completeness, here is a modified proof.

Proof. From a factorization of $f: C \rightarrow D$ through a non-trivial unramified map $\tilde{D} \rightarrow D$, we get a partition of $f^{-1}(\star)$ according to the different images in \tilde{D} . The monodromy of f has to preserve this partition, and hence cannot be the full symmetric group.

Suppose now that f is primitive. We want to show it has full monodromy. Let K be the kernel of the map $i_*: \pi_1(D^\times, \star) \rightarrow \pi_1(D, \star)$. We start by showing the following.

Claim 4.2. *The action of K on $f^{-1}(\star)$ is transitive.*

Proof. Consider the following commutative diagram of groups:

$$\begin{array}{ccccc} \pi_1(C^\times, \star) & \xrightarrow{i_*} & \pi_1(C, \star) & & \\ \downarrow f_* & & \downarrow f_* & & \\ K \longrightarrow \pi_1(D^\times, \star) & \xrightarrow{i_*} & \pi_1(D, \star) & & \end{array}$$

Given any element $g \in \pi_1(D^\times, \star)$, we can use the surjectivity of the top and right arrows to find an element $\bar{g} \in \pi_1(C^\times, \star)$ such that

$$i_*(g) = f_* i_*(\bar{g}) = i_*(f_*(\bar{g}))$$

Hence,

$$g \cdot (f_*(\bar{g}))^{-1} = \kappa \in K$$

Now let us consider the monodromy action of g on the fiber $f^{-1}(\star)$. More specifically, let us see how g acts on the base point $\star \in C^\times$ (which in turn lies over the base point $\star \in D^\times$).

If we lift a representative loop of $f_*(\bar{g})$ with start point \star , we will get a representative loop of \bar{g} . Hence, the monodromy action of $f_*(\bar{g})$ fixes \star . As $g = \kappa \cdot f_*(\bar{g})$, we get $g(\star) = \kappa(\star)$.

We showed that for any $g \in \pi_1(D^\times, \star)$, there is a $\kappa \in K$ that acts on \star in the same way. Hence, the orbit of \star under the K -action is equal to the orbit under the $\pi_1(D^\times, \star)$ -action. But the latter action is transitive, since C^\times is connected. Hence, the action of K is transitive as well. \square

Claim 4.3. *The image of $K \subset \pi_1(D^\times, \star) \rightarrow S_d$ is generated by transpositions.*

Proof. We use the standard presentations of $\pi_1(D, \star)$ and $\pi_1(D^\times, \star)$. Let $g(D)$ be the genus of D . Then $\pi_1(D, \star)$ is the free group generated by α_i, β_i for i ranging from 1 to $g(D)$, subject to the relation

$$[\alpha_1, \beta_1] \dots [\alpha_{g(D)}, \beta_{g(D)}] = \text{Id}$$

Let b be the number of branch points of $f: C \rightarrow D$, and choose simple contractible loops τ_j around each branch point. Then $\pi_1(D^\times, \star)$ is the free group generated by α_i, β_i for $i = 1, 2, \dots, g(D)$ and τ_j for $j = 1, 2, \dots, b$, subject to the relation

$$[\alpha_1, \beta_1] \dots [\alpha_{g(D)}, \beta_{g(D)}] = \tau_1 \tau_2 \dots \tau_b$$

The map $\pi_1(D^\times, \star) \rightarrow \pi_1(D, \star)$ sends the α_i, β_i to the corresponding α_i, β_i , and each τ_j to Id . Hence, the kernel K is the normal closure of the subgroup generated by the τ_j .

Since $f: C \rightarrow D$ is simply branched, each τ_j acts as a transposition. Hence so will any conjugate of τ_j . Therefore, the image of K in S_d is generated by transpositions, as we wanted to show. \square

Now we can finish the proof of Proposition 4.1. The image of K in S_d is transitive and generated by transpositions. But the only subgroup of S_d satisfying these two properties is S_d itself. Hence, $K \rightarrow S_d$ is surjective, and so is the monodromy map $\varphi: \pi_1(D^\times, \star) \rightarrow S_d$. \square

Corollary 4.4. *A primitive degree d simply branched cover $f: C \rightarrow D$ does not admit any isomorphisms over D , unless $d = 2$.*


Proof. This is similar to Proposition 1.9 in [Ful69]. An isomorphism of C over D corresponds to an element σ of S_d that commutes with every element of the monodromy group $\text{im } \varphi \subset S_d$.

By Proposition 4.6, if a cover is primitive, then the monodromy map is surjective, and therefore σ would be in the center of S_d , which is the trivial group as soon as $d \geq 3$. \square

An arbitrary cover may be related to primitive ones by the following construction.

Corollary 4.5. *Let $f: C \rightarrow D$ be a simply branched covering map. Then it can be factored in a unique way as $C \rightarrow \bar{D} \rightarrow D$ such that $C \rightarrow \bar{D}$ has full monodromy, and $\bar{D} \rightarrow D$ is unramified.*

Proof. The unramified cover $\bar{D} \rightarrow D$ is the covering space corresponding to the subgroup $\text{im } \pi_1(C, \star) \subset \pi_1(D, \star)$. By the lifting property, the map $C \rightarrow D$ factors

FIGURE 9. The preimages of \star in \overline{D} and C .

through \overline{D} . The map $C \rightarrow \overline{D}$ is surjective on π_1 . Hence, by Proposition 4.1, it has full monodromy. \square

Corollary 4.5 gives us a tool to study the monodromy group of an arbitrary simply branched cover $f: C \rightarrow D$. We would hope to relate the monodromy group of f with the monodromy groups of each factor, which are easier to understand. The following proposition is a first step in this direction.

Proposition 4.6. *Let $C \rightarrow \overline{D} \rightarrow D$ be as in Corollary 4.5. Let G (resp. \overline{G}) be the monodromy group of $C \rightarrow D$ (resp. $\overline{D} \rightarrow D$). The following holds.*

- (1) *There is a surjective map $G \twoheadrightarrow \overline{G}$,*
- (2) *There kernel of $G \twoheadrightarrow \overline{G}$ is as large as it can be. That is, the monodromy group G fits in a exact sequence*


$$1 \rightarrow S_{\tilde{d}} \times \dots \times S_{\tilde{d}} \rightarrow G \rightarrow \overline{G} \rightarrow 1$$

where the product in the left term has $\tilde{d} = \deg(\overline{D} \rightarrow D)$ factors, each of them isomorphic to the symmetric group on $\tilde{d} = \deg(C \rightarrow \overline{D})$ elements.

Proof. We will define a map $G \rightarrow \overline{G}$ such that the monodromy maps are compatible. That is, we want the following diagram to commute

$$\begin{array}{ccc} \pi_1(D^\times, \star) & \longrightarrow & G \\ \downarrow i_* & & \downarrow \varphi \\ \pi_1(D, \star) & \xrightarrow{\overline{\varphi}} & \overline{G} \end{array}$$

The map is defined as follows. Given a monodromy element $g \in G$, choose a representative loop $\gamma: I \rightarrow D^\times$ inducing it. We send g to $\overline{g} = \overline{\varphi}(i_*(\gamma))$. Geometrically, \overline{g} is the monodromy element induced by the intermediate lifts of γ to \overline{D} :



The map $G \rightarrow \overline{G}$ is well defined because if a loop $[\gamma] \in \pi_1(D^\times, \star)$ induces the trivial monodromy on the fibers of $C \rightarrow D$, it also induces the trivial monodromy on the fibers of the intermediate map $\overline{D} \rightarrow D$.

Since $i_*: \pi_1(D^\times, \star) \rightarrow \pi_1(D, \star)$ and $\overline{\varphi}: \pi_1(D, \star) \rightarrow \overline{G}$ are surjective, so is the map $G \rightarrow \overline{G}$. This establishes item 1.

Let us introduce some notation. Let $\bar{f}: C \rightarrow \overline{D}$ and $h: \overline{D} \rightarrow D$ be the factors of the factorization of $f: C \rightarrow D$. Call $p_1, \dots, p_{\bar{a}} \in \overline{D}$ the points in the preimage of $\star \in D$, and let $S_i = \bar{f}^{-1}(p_i) \subset C$. See fig. 9.

Let K be the kernel of $G \twoheadrightarrow \overline{G}$. The elements of K preserve the sets S_i for all i . That is, $K \subset \text{Sym}(S_1) \times \text{Sym}(S_2) \times \dots \times \text{Sym}(S_{\bar{a}}) \subset \text{Sym}(\cup S_i)$. We want to show that $K = \text{Sym}(S_1) \times \text{Sym}(S_2) \times \dots \times \text{Sym}(S_{\bar{a}})$.

Pick a simple loop $\gamma \in \pi_1(D^\times, \star)$, contractible in D , around one of the branch points. Let $\tau = \varphi(\gamma) \in G$ be the corresponding monodromy element. It is a transposition, since the branching is simple. Moreover, since γ is contractible in D , the transposition τ maps to the identity in \overline{G} . Hence $\tau \in K$. The transposition τ has to act trivially in all but one of the S_i . Say it acts non-trivially on S_1 . By a slight abuse of notation, we will say that $\tau \in \text{Sym}(S_1)$.

We will exploit the action of G on K by conjugation. First, consider the composition of the pushforward $h_*: \pi_1(\overline{D}^\times, p_1) \rightarrow \pi_1(D^\times, \star)$ with the monodromy map $\varphi_f: \pi_1(D^\times, \star) \rightarrow G$. The composite map $\varphi_f \circ h_*$ fixes S_1 , and hence it factors through $\text{Sym}(S_1) \times \text{Sym}(S_2 \cup \dots \cup S_{\bar{a}})$. Projecting on the $\text{Sym}(S_1)$ factor, we recover

the monodromy map of $C \rightarrow \overline{D}$, as in the following diagram.

$$\begin{array}{ccc}
 & & \text{Sym}(S_1) \\
 & \nearrow \varphi_{\overline{f}} & \uparrow \\
 \pi_1(\overline{D}^\times, p_1) & \dashrightarrow & \text{Sym}(S_1) \times \text{Sym}(S_2 \cup \dots \cup S_{\overline{d}}) \\
 \downarrow h_* & & \downarrow \\
 \pi_1(D^\times, \star) & \xrightarrow{\varphi_f} & G \subset \text{Sym}(S_1 \cup S_2 \cup \dots \cup S_{\overline{d}})
 \end{array}$$

Pick $\gamma \in \pi_1(\overline{D}^\times, p_1)$. Conjugating an element of $\text{Sym}(S_1) \cap K$ by $\varphi_f(h_*(\gamma))$ is the same as conjugating by $\varphi_{\overline{f}}(\gamma) \in \text{Sym}(S_1)$. As $C \rightarrow \overline{D}$ has full monodromy, we can get an arbitrary conjugate in $\text{Sym}(S_1)$.

We already have a transposition $\tau \in \text{Sym}(S_1) \cap K$. Hence, we can get *all* transpositions in $\text{Sym}(S_1)$. But these generate the full symmetric group. Hence, $\text{Sym}(S_1) \times \text{Id} \times \text{Id} \times \dots \times \text{Id} \subset K$.

For each i , pick $\overline{g} \in \overline{G}$ that sends p_1 to p_i . Lift \overline{g} to $g \in G$. Then

$$g \text{Sym}(S_1) g^{-1} = \text{Sym}(S_i) \subset K$$

As this holds for any i , we get

$$\text{Sym}(S_1) \times \text{Sym}(S_2) \times \dots \times \text{Sym}(S_{\overline{d}}) \subset K$$

as we wanted to show. \square

Note that the exact sequence in Proposition 4.6 does not completely characterize the monodromy group G —we do not know what extension it is. Determining what extensions do arise as monodromy groups of simply branched maps is an interesting question, which, as far as I know, is unsolved.

We now apply Proposition 4.6 to establish the result we will need in Section 5.

Lemma 4.7. *Let C be a smooth curve, $C \rightarrow \overline{D}$ be a simply branched map with full monodromy, and $\overline{D} \rightarrow D$ an unramified map. Consider the map*

$$C \times_D C - \Delta \rightarrow \overline{D} \times_D \overline{D}$$

The target $\overline{D} \times_D \overline{D}$ is smooth, but possibly disconnected. Nevertheless, the preimage of each component is irreducible.

Proof. Let $\Gamma \subset \overline{D} \times_D \overline{D}$ be a component. We want to show that its preimage is connected.

If $\Gamma = \Delta$ is the diagonal, having connected preimage is equivalent to the monodromy group of $C \rightarrow \overline{D}$ being 2-transitive. This is immediate, since the monodromy group is actually the full symmetric group.

Suppose now Γ is not the diagonal. For convenience, let us use the notation as in the proof of Proposition 4.1. Pick a point in the fiber over \star of the map $\Gamma \rightarrow D$, say (p_i, p_j) . For any $q_i, q'_i \in S_i$ and $q_j, q'_j \in S_j$, we want to exhibit a path in $C \times_D C$ connecting (q_i, q_j) to (q'_i, q'_j) . It is enough to find an element in the monodromy group of $C \rightarrow D$ that sends q_i to q'_i , and q_j to q'_j . This follows from Proposition 4.1, since the monodromy group contains $\text{Sym}(S_i) \times \text{Sym}(S_j)$. \square

5. IRREDUCIBILITY OF THE HURWITZ SPACE $\mathcal{H}_{N,g}^0(E)$

Let us resume our study of Hurwitz spaces. Recall their definition.

Definition 5.1. The Hurwitz space $\mathcal{H}_{N,g}(E)$ parametrizes simply branched covers (meaning, dominant finite flat morphisms) $C \rightarrow E$ from a smooth genus g curve C to the fixed curve E . The space $\mathcal{H}_{N,g}^0(E)$ is the open and closed subscheme of $\mathcal{H}_{N,g}(E)$ of primitive covers $C \rightarrow E$ (that is, covers which do not factor through any non-trivial unramified maps $\tilde{E} \rightarrow E$).

There are several alternate characterizations of $\mathcal{H}_{N,g}^0(E)$. For example, as a consequence of Proposition 4.1 and Corollary 4.5, we have the following useful one.

Proposition 5.2. *The space $\mathcal{H}_{N,g}^0(E) \subset \mathcal{H}_{N,g}(E)$ consists of the covers $C \rightarrow E$ with full monodromy.*

Our goal is to prove Theorem 1.4, which says that for a smooth genus one curve E and $g > 1$, the Hurwitz space $\mathcal{H}_{N,g}^0(E)$ is irreducible. Note that from the irreducibility of $\mathcal{H}_{N,g}^0(E)$, we can characterize all the irreducible components of $\mathcal{H}_{N,g}(E)$.

Corollary 5.3. *The components of $\mathcal{H}_{N,g}(E)$ are in bijection with isogenies $\tilde{E} \rightarrow E$ of degree dividing N . That is, the decomposition in irreducible components is*

$$\mathcal{H}_{N,g}(E) = \bigsqcup_{\tilde{E} \rightarrow E} \mathcal{H}_{N',g}^0(\tilde{E})$$

where $N' = N / \deg(\tilde{E} \rightarrow E)$, for $N' \geq 2$.

Proof. The right hand side is included in the left by the map

$$(C \rightarrow \tilde{E}) \mapsto (C \rightarrow \tilde{E} \rightarrow E)$$

Corollary 4.5 says that the opposite inclusion holds, and Theorem 1.4 says every term in the right is irreducible. \square

5.1. Our strategy. As a motivating example, let us go over a simple argument by Fulton for the irreducibility of \mathcal{M}_g (in the appendix of [HM82]). We follow Harris and Morisson account in [HM98].

The idea, which really goes back to Deligne and Mumford, is to look instead at the compactified $\overline{\mathcal{M}}_g$ and leverage the structure of the boundary to prove irreducibility. More precisely, assume by induction that $\mathcal{M}_{g'}$ is irreducible for $g' < g$. Then, by induction, one can establish that each boundary divisor Δ_i in $\overline{\mathcal{M}}_g$ is irreducible. Moreover, one can prove that any pair of boundary divisors meet, simply by exhibiting an element in the intersection $\Delta_i \cap \Delta_j$. Hence, there is a unique connected component of $\overline{\mathcal{M}}_g$ meeting the boundary Δ .

However, $\overline{\mathcal{M}}_g$ is locally irreducible, and therefore connectedness implies irreducibility. Hence, we established the following reduction.

Reduction. *It is enough to show that every component Γ of $\overline{\mathcal{M}}_g$ meets the boundary Δ .*

To exhibit such a curve, we will pass to a more structured space: the Hurwitz space $\mathcal{H}_{d,g}(\mathbb{P}^1)$ of degree d covers $C \rightarrow \mathbb{P}^1$. It follows from Riemann–Roch that for large d the Hurwitz space dominates \mathcal{M}_g .

If we were willing to use that the Hurwitz space is irreducible, we would conclude irreducibility of \mathcal{M}_g as well. However, using the reduction above, we can still prove irreducibility of $\overline{\mathcal{M}}_g$ without invoking the irreducibility of $\mathcal{H}_{d,g}(\mathbb{P}^1)$. Instead, we will use the properness of the admissible covers compactification $\overline{\mathcal{H}}_{d,g}(\mathbb{P}^1)$.

The branch morphism extends to the boundary of the admissible covers compactification, and every component of $\overline{\mathcal{H}}_{d,g}(\mathbb{P}^1)$ dominates the Fulton–MacPherson configuration space $\mathbb{P}^1[b]$. By choosing an extremal configuration of branch points, we can force the cover to be a reducible curve whose all irreducible components are rational. In particular, it is an element of Δ . For details, see [HM98, Chapter 6A].

We will apply a similar argument to show the irreducibility of $\mathcal{H}_{N,g}^0(E)$. Again, the thrust of the argument comes from exploiting boundary divisors in a compactification of $\mathcal{H}_{N,g}^0(E)$. So let us start by choosing one.

A natural candidate is to use the Kontsevich space $\overline{\mathcal{M}}_g(E, N)$. However, this can have bad singularities at the locus of stable maps with contracted components, and it is not clear that it is locally irreducible, which is key to an argument like the one above.

We will instead use a partial compactification which is smooth, but still allows enough singular covers $X \rightarrow E$ for our reduction argument to go through.

Definition 5.4. Let $\tilde{\mathcal{H}}_{N,g}(E) \subset \overline{\mathcal{M}}_g(E, N)$ to be the open substack parametrizing finite maps (that is, the complement of the locus where there are contracted components). Points in $\tilde{\mathcal{H}}_{N,g}(E)$ correspond to degree N , dominant, finite flat maps from a connected, arithmetic genus g nodal curve X to E . There is an open and closed substack $\tilde{\mathcal{H}}_{N,g}^0(E) \subset \tilde{\mathcal{H}}_{N,g}(E)$ parametrizing primitive covers.

Remark 5.5. Not being primitive is naturally a closed property, since if a family of covers $X_t \rightarrow E$ factors through $\tilde{E} \rightarrow E$ for generic t , then it factors through for all t . This implies that $\tilde{\mathcal{H}}_{N,g}^0(E) \subset \tilde{\mathcal{H}}_{N,g}(E)$ is an open substack.

However, it is also closed. Indeed, say $\tilde{E} \rightarrow E$ is a non-trivial isogeny, and the composition $X_0 \rightarrow \tilde{E} \rightarrow E$ is a stable map which is a limit of primitive maps $X_t \rightarrow E$. Then $(X_0 \rightarrow E)$ is a singular point of $\overline{\mathcal{M}}_g(E, N)$. Indeed, it lies in at least two components: the closure of $\mathcal{H}_{N,g}^0(E)$, and the closure of $\mathcal{H}_{N',g}^0(\tilde{E})$, where N' is the degree of $X_0 \rightarrow \tilde{E}$. Both of these components are already $2g - 2$ dimensional, so $(X_0 \rightarrow E)$ is a singular point. But, as we will see below in Lemma 5.6, $\tilde{\mathcal{H}}_{N,g}(E)$ lives in the smooth locus of $\overline{\mathcal{M}}_g(E, N)$, so we would not see such a $X_0 \rightarrow E$.

Lemma 5.6. *The space $\tilde{\mathcal{H}}_{N,g}(E)$ is non-singular.*

Proof. The tangent space at a point $f : C \rightarrow E$ is equal to $H^0(C, \mathcal{N}_f)$, where the normal sheaf \mathcal{N}_f is defined as

$$0 \rightarrow T_C \rightarrow f^*T_E \rightarrow \mathcal{N}_f \rightarrow 0$$

When there are no contracted components, \mathcal{N}_f is torsion. Hence, $H^1(C, \mathcal{N}_f) = 0$, and Riemann–Roch says $h^0(C, \mathcal{N}_f) = 2g - 2$, the length of the branch subscheme. As $\dim \tilde{\mathcal{H}}_{N,g}(E)$ is $2g - 2$ as well, we conclude it is non-singular. \square

We will show that $\tilde{\mathcal{H}}_{N,g}^0(E)$ is connected. Of course, Theorem 1.4 follows from this. Our proof will be inductive.

Induction Hypothesis 1. Suppose, by induction, that the Hurwitz space of primitive covers $\tilde{\mathcal{H}}_{N,g'}^0(E)$ is connected for all genus one curves E , and integers $N \geq 2$ and $2 \leq g' < g$.

The base case $g = 2$ is due to Kani in [Kan03].

Theorem (Kani). *The Hurwitz space $\mathcal{H}_{N,2}^0(E)$ is irreducible.*

We will assume $g \geq 3$ from now on.

As in Fulton's argument, let us take advantage of the boundary in our partial compactification. Call $\Delta \subset \tilde{\mathcal{H}}_{N,g}^0(E)$ the locus of covers with singular source.

As opposed to the case of \mathcal{M}_g , where there are relatively few boundary divisors and they are easily understood by induction, our boundary Δ will have quite a few boundary components, and they are not so easily dealt with inductively. Nevertheless, we will prove the following.

Final Reduction 1. *It is enough to show that every connected component of $\tilde{\mathcal{H}}_{N,g}^0(E)$ contains a point in the boundary Δ . That is, it is enough to find, in each component, a cover with singular source.*

In Section 5.2 we will prove Final Reduction 1 under the Induction Hypothesis 1. Then in Section 5.3 we will establish that every component of Hurwitz space does contain a point of Δ . As in Fulton's argument, we have to pass to a more structured space to find such a point. In our case, we will use the Severi variety of curves in $E \times \mathbb{P}^1$. Specifically, we will invoke Theorem 1.20.

5.2. The reduction argument. In this section we will prove Final Reduction 1 under the Induction Hypothesis 1. To do so, we will first focus on a few special boundary divisors, and then bootstrap to the whole boundary Δ .

We define $D_{\text{sing}} \subset \tilde{\mathcal{H}}_{N,g}^0(E)$ to be the closure of the locus of covers $X \rightarrow E$ with X irreducible and with a single node, for which the composition of the normalization map $\tilde{X} \rightarrow X$ and $X \rightarrow E$ is primitive. That is, $\tilde{X} \rightarrow E$ does not factor through non-trivial isogenies, as in the diagram in fig. 10.


Lemma 5.7. *Under Induction Hypothesis 1, the set D_{sing} is an irreducible divisor.*

Proof. We may parametrize the points in D_{sing} by the data of a primitive cover $\pi: C \rightarrow E$ in $\mathcal{H}_{d,g-1}^0(E)$, and a pair of distinct points $p, q \in C$ in a fiber of π . This data is recorded in the fiber product $\Sigma = \mathcal{C} \times_E \mathcal{C} - \Delta$, where \mathcal{C} is the universal curve over $\mathcal{H}_{N,g-1}^0(E)$, and Δ is the diagonal. There is a birational isomorphism from Σ to D_{sing} that identifies the points p and q . Note that

$$\dim D_{\text{sing}} = \dim \Sigma = 2(g-1) - 2 + 1 = 2g - 3 = \dim \mathcal{H}_{d,g}^0(E) - 1$$

Hence D_{sing} is a divisor. Moreover, from the induction hypothesis, we know that $\mathcal{H}_{N,g-1}^0(E)$ is irreducible. From Proposition 5.2, the monodromy of $C \rightarrow E$ is full, and therefore 2-transitive. Hence, Σ is irreducible, and so is D_{sing} . \square

From the Lemma 5.7, we get our first reduction.

FIGURE 10. General point of D_{sing}

Reduction 1. *It is enough to show that every connected component of $\tilde{\mathcal{H}}_{N,g}^0(E)$ contains a point in D_{sing} .*

Proof. By Lemma 5.6, $\tilde{\mathcal{H}}_{N,g}^0(E)$ is smooth, and hence it is enough to show it is connected. However, by Lemma 5.7, only one connected component meets D_{sing} . \square

Let us define a few more special boundary divisors. For each isogeny $\pi: \tilde{E} \rightarrow E$, let $D_\pi \subset \tilde{\mathcal{H}}_{N,g}^0(E)$ be the closure of the subscheme parametrizing covers with reducible source $X = X' \cup_p \tilde{E}$, with a single node p , such that $X' \rightarrow E$ is a primitive cover of degree $N - \deg \pi$, and genus $g - 1$. See fig. 11 for a picture of the general point of D_π .

Remark 5.8. Given two pointed covers $f_i: (C_i, p_i) \rightarrow E$, for $i = 1, 2$, we can form a cover $f: C \rightarrow E$ by attaching C_2 to C_1 in the following way. Take


$$C = C_1 \sqcup C_2 / (p_1 \sim p_2)$$

and glue f_1 to $\tau \circ f_2$ where $\tau: E \rightarrow E$ is the translation which takes $f_2(p_2)$ to $f_1(p_1)$.

We may omit the point $p_i \in C$ in case $C_i \rightarrow E$ is an isogeny, since then the choice of point is irrelevant.

For example, a cover in D_π is obtained by attaching the isogeny π to a cover $X' \rightarrow E$ in $\mathcal{H}_{N-\deg \pi, g-1}^0(E)$.

Lemma 5.9. *Under Induction Hypothesis 1, the set D_π is an irreducible divisor.*

FIGURE 11. General point of D_π

Proof. There is a map from the universal curve $\mathcal{C} \rightarrow \tilde{\mathcal{H}}_{N-\deg \pi, g-1}^0(E)$ to D_π , which takes a pointed cover $(X', p) \rightarrow E$ and attaches the isogeny $\pi: \tilde{E} \rightarrow E$ at the marked point p . We conclude that D_π is a divisor, since

$$\begin{aligned} \dim D_\pi &= \dim \mathcal{C} = \dim \mathcal{H}_{N-\deg \pi, g-1}^0 + 1 \\ &= 2(g-1) - 2 + 1 = 2g - 3 = \dim \mathcal{H}_{N, g}^0(E) - 1 \end{aligned}$$

Moreover, by induction, the universal curve \mathcal{C} is irreducible, and so is D_π . \square

We are on our way to prove

Reduction 2. *It is enough to show that every connected component of $\tilde{\mathcal{H}}_{N, g}^0(E)$ contains a point in D_{sing} or in D_π for some isogeny π .*

Indeed, this follows will follow from the Reduction 1 and the following proposition.


Proposition 5.10. *The union*

$$D_{\text{sing}} \cup \bigcup_{\pi} D_\pi \subset \tilde{\mathcal{H}}_{N, g}^0(E)$$

is connected, where π ranges through all isogenies $\pi: \tilde{E} \rightarrow E$ of degree at most $N - 1$.

In view of irreducibility of D_π and D_{sing} , it is enough to prove that these divisors meet for every π . That is, we only have to exhibit a curve that is in both.

Here is a candidate. Let d be the degree of the isogeny $\pi: \tilde{E} \rightarrow E$. Take a cover $\tilde{X}_0 \rightarrow E$ in $\mathcal{H}_{N-d, g-2}(E)$. Now pick two points $p_1, p_2 \in \tilde{X}_0$ in a fiber of the map

FIGURE 12. A cover in both D_π and D_{sing}

to E , and glue them to get a nodal curve $X_0 \rightarrow E$. Then, pick a generic third point $q \in X_0$, and attach a copy of \tilde{E} at q . Let this be $X \rightarrow E$, as in fig. 12.

From the definitions of D_π and D_{sing} , we make the following observation.

Observation 5.11. If $\tilde{X}_0 \cup \tilde{E} \rightarrow E$ and $X \rightarrow E$ are both primitive covers, then $X \rightarrow E$ is a cover in the intersection $D_{\text{sing}} \cap D_\pi$.


Corollary 5.12. *If $g \geq 4$, then D_{sing} meets D_π for every isogeny π .*

Proof. When $g - 2 \geq 2$, we may choose a primitive cover $\tilde{X}_0 \rightarrow E$. This guarantees that both conditions in Observation 5.11 are satisfied. \square

This settles Proposition 5.10 for $g \geq 4$, and in consequence the Reduction 2.

We can apply the same construction for $g = 3$, but then $\tilde{X}_0 \rightarrow E$ would be an isogeny. We have to choose it carefully so that the hypotheses of Observation 5.11 are satisfied. It turns out that we cannot choose such a $\tilde{X}_0 \rightarrow E$ for every isogeny π , but we can for a large subclass: the reduced ones.

Definition 5.13. An isogeny $\pi: \tilde{E} \rightarrow E$ is *reduced* if it does not factor through the multiplication by n map $M_n: E \rightarrow E$ for any integer $n > 1$.

FIGURE 13. A cover in both D_{π_1} and D_{π_2} .

Lemma 5.14. *For any reduced isogeny $\pi: \tilde{E} \rightarrow E$ and any $D \geq 2$, there is a reduced degree D isogeny $\hat{\pi}: \hat{E} \rightarrow E$ and points $p_1, p_2 \in \hat{E}$ such that*

- (1) *There is no common non-trivial subcover to π and $\hat{\pi}$.*
- (2) *For any subcover $\pi': \hat{E} \rightarrow E'$, the points $\pi'(p_1), \pi'(p_2) \in E'$ are the same if and only if $\pi' = \hat{\pi}$ (that is, the subcover is trivial).*

We will prove Lemma 5.14 in Section 5.4.

Corollary 5.15. *For $g = 3$ and a reduced isogeny π , the divisor D_π meets D_{sing} .*

Proof. Take \tilde{X}_0 as the \hat{E} produced by Lemma 5.14. This satisfies the conditions of Observation 5.11. \square

To finish the proof of Proposition 5.10, for each isogeny π , we need to find a reduced isogeny π' such that D_π meets $D_{\pi'}$.

Lemma 5.16. *For $g = 3$, the divisors D_{π_1} and D_{π_2} meet if $d_1 + d_2 = N - 1$, where d_i is the degree of the isogeny $\pi_i: \tilde{E}_i \rightarrow E$.*

Proof. Form the nodal curve X with components \tilde{E}_1, \tilde{E}_2 and E , by gluing \tilde{E}_1 with E at a point p , and \tilde{E}_2 with E at a point $q \neq p$, as in fig. 13. The cover $X \rightarrow E$ lies in both D_{π_1} and D_{π_2} . \square

This completes the proof of Proposition 5.10, since there are reduced isogenies of any degree. This is easy to see directly, but is already a consequence of Lemma 5.14, so we are not going to restate it as a separate claim.

We have so far settled Reduction 2. This gives a good mileage towards our goal, the Final Reduction 1. To help organize what is left to do, let us introduce a piece of language.

Definition 5.17. Let $C \rightarrow E$ be a cover in $\tilde{\mathcal{H}}_{N,g}^0(E)$, and $\eta \in C$ a node in the source. Let \tilde{C} be the partial normalization of C at η . We call η

- a *disconnecting node* if \tilde{C} is disconnected.
- a *E-superfluous node* if \tilde{C} is connected and $\tilde{C} \rightarrow E$ is primitive.
- a *E-relevant node* if $\tilde{C} \rightarrow E$ is not primitive, i.e., it factors through a non-trivial isogeny.

In this language, a cover is in D_{sing} if and only if it has an E-superfluous node. We can reexpress Reduction 2 as saying that it is enough to find, in each connected component, a cover with an E-superfluous node, or a disconnecting node such that if the two components of the partial normalization are $C_1 \sqcup C_2$, then

- (1) $C_1 \rightarrow E$ is an isogeny, and
- (2) $C_2 \rightarrow E$ is primitive.

We can relax this and drop condition (1).

Reduction 3. *It is enough to find, in each connected component of Hurwitz space $\mathcal{H}_{N,g}^0(E)$, a cover $C \rightarrow E$ with an E-superfluous node or a disconnecting node such that the restriction of $C_1 \sqcup C_2 \rightarrow C \rightarrow E$ to $C_2 \rightarrow E$ is primitive.*

Proof. Indeed, let us say we have a cover $C \rightarrow E$ and a disconnecting node $\eta \in C$ such that $C_2 \rightarrow E$ is primitive. If $C_1 \rightarrow E$ is an isogeny, this is a point in D_π , and we are done. Assume otherwise that $C_1 \rightarrow E$ is not an isogeny. We want to find another point in the same connected component which is in one of the D_π divisors.

Let $\pi: \tilde{E} \rightarrow E$ the maximal isogeny for which $C_1 \rightarrow E$ factors through, and let $\eta_1 \in C_1$ be the point over the node η . Then $(C_1, \eta_1) \rightarrow \tilde{E}$ is an element of $\mathcal{C} \rightarrow \tilde{\mathcal{H}}_{N_1,g_1}^0(\tilde{E})$, the universal curve over the component of Hurwitz space $C_1 \rightarrow E$ lies on. By the Induction Hypothesis 1, \mathcal{C} is irreducible. Hence, we can connect (C_1, η_1) to any other element of \mathcal{C} . For example, we can take $(C' \cup_p \tilde{E}, \eta') \rightarrow \tilde{E}$, where $\eta' \in C'$, for some cover $C' \rightarrow \tilde{E}$.

Attaching $(C_2, \eta) \rightarrow E$ back to $(C' \cup_p \tilde{E}, \eta') \rightarrow \tilde{E} \rightarrow E$, we get

$$X = \tilde{E} \cup_p C' \cup_\eta C_2 \rightarrow E$$

as in fig. 14. As $C_2 \rightarrow E$ is primitive, so is $C' \cup_\eta C_2 \rightarrow E$. Hence, the cover $X \rightarrow E$ is an element of D_π . \square

Remark 5.18. Note that for this argument it was important that our induction hypothesis accommodates for any genus one target, as long as the genus of the source is smaller than g .

We can do even better.

Reduction 4. *It is enough to find, in each connected component of Hurwitz space, a cover with an E-superfluous node or a disconnecting node.*




FIGURE 14. From a disconnecting node with a primitive component to a cover in D_{π} , as in Reduction 3.





FIGURE 15. From a cover with a disconnecting node to a cover with a disconnecting node and a primitive component, as in Reduction 4.

Proof. We use the same trick. Let $C_1 \sqcup C_2 \rightarrow C \rightarrow E$ be the partial normalization of a curve with a disconnecting node η . Let $\eta_i \in C_i$ be the point over η . As $g \geq 3$, at least one of the C_i has genus at least 2. Say that is C_1 . Let $\tilde{E} \rightarrow E$ be the maximal isogeny $C_1 \rightarrow E$ factors through. Degenerate $(C_1, \eta_1) \rightarrow E$ to $(C' \cup \tilde{E}, \tilde{\eta}) \rightarrow E$ with $\tilde{\eta} \in \tilde{E}$. Attaching (C_2, η_2) back, we get $C' \cup_p \tilde{E} \cup_{\eta} C_2 \rightarrow E$, as in fig. 15.

We claim that cover $\tilde{E} \cup_{\eta} C_2 \rightarrow E$ is primitive. Suppose not. Then, by definition, the map $\tilde{E} \cup_{\eta} C_2 \rightarrow E$ factors through a non-trivial isogeny, and so would the original cover $C_1 \cup_{\eta} C_2 \rightarrow E$ (because $C_1 \rightarrow E$ itself factors through $\tilde{E} \rightarrow E$). But the original cover is primitive by assumption, which is a contradiction.

Hence, the node p joining C' and \tilde{E} is a disconnecting node for which one of the components is primitive. So we are done by Reduction 3. \square

Let us finally prove the Final Reduction 1.

FIGURE 16. An E-relevant node η .

Proof. Let $C \rightarrow E$ be a cover in $\Delta \subset \tilde{\mathcal{H}}_{N,g}^0(E)$. Hence, C has a node, call it η . By Reduction 4, it is enough to show that in the connected component of $C \rightarrow E$ there is a cover with an E-superfluous or a disconnecting node.

If η itself is E-superfluous or disconnecting, we are done. We may assume that η is E-relevant, then. This means that the partial normalization $\tilde{C} \rightarrow E$ is connected, and that it factors through a non-trivial isogeny $\tilde{E} \rightarrow E$. Let $\eta_1, \eta_2 \in \tilde{C}$ be the preimages of the node η , and $N' = \deg(\tilde{C} \rightarrow \tilde{E})$, as in fig. 16.

The data of $C \rightarrow E$ is encoded in the cover $\tilde{C} \rightarrow \tilde{E}$, with the two marked points η_1, η_2 . Let us construct a parameter space \mathcal{P} for the triples $(\tilde{C} \rightarrow \tilde{E}, \eta_1, \eta_2)$.

Let $\mathcal{C} \rightarrow \tilde{\mathcal{H}}_{N',g-1}^0(\tilde{E})$ be the universal curve over the component of Hurwitz parametrizing primitive covers. The individual covers $C_b \rightarrow \tilde{E}$ assemble together in a map $\mathcal{C} \rightarrow \tilde{E}$. Then the triple $(\tilde{C} \rightarrow \tilde{E}, \eta_1, \eta_2)$ is a point in


$$\mathcal{P} = \mathcal{C} \times_{E \times \tilde{\mathcal{H}}} \mathcal{C} - \Delta$$

where the map from \mathcal{C} to E is the composition of $\mathcal{C} \rightarrow \tilde{E}$ with the isogeny $\tilde{E} \rightarrow E$.

Conversely, given any point $(\tilde{C}' \rightarrow \tilde{E}, \tilde{\eta}_1, \tilde{\eta}_2) \in \mathcal{P}$, we may identify the points $\tilde{\eta}_1$ and $\tilde{\eta}_2$ to get a cover $C' \rightarrow E$ in $\tilde{\mathcal{H}}_{N,g}^0(E)$.

This prompts us to characterize the connected component of \mathcal{P} containing the point $(\tilde{C} \rightarrow \tilde{E}, \eta_1, \eta_2)$. If this connected component contains a point which gives rise to a cover of E with a disconnecting or E-superfluous node, we will be done.

Let κ_i be the image of η_i in \tilde{E} . If $\kappa_1 = \kappa_2$, then $C \rightarrow E$ would also factor through $\tilde{E} \rightarrow E$, which is contrary to our assumption that $C \rightarrow E$ is primitive. Hence $\kappa_1 \neq \kappa_2$.

FIGURE 17. The construction for the case $g \geq 4$.

We can even say more: after identifying κ_1 with κ_2 , the map from the nodal curve $\tilde{E}/(\kappa_1 \sim \kappa_2)$ to E is a primitive cover.

Note that the evaluation map $\mathcal{C} \rightarrow \tilde{E}$ produces a map

$$f: \mathcal{P} = \mathcal{C} \times_{E \times \tilde{E}} \mathcal{C} - \Delta \rightarrow \tilde{E} \times_E \tilde{E}$$

Our point $(\tilde{C} \rightarrow \tilde{E}, \eta_1, \eta_2)$ in \mathcal{P} maps to (κ_1, κ_2) .

We describe the components of \mathcal{P} in terms of the map f .

Proposition 5.19. *The map $f: \mathcal{P} \rightarrow \tilde{E} \times_E \tilde{E}$ induces a bijection in connected components.*


Before we prove this proposition, let us see how to establish the Final Reduction 1 assuming it.

If $g \geq 4$, we will use the construction in Figure 17. Take a primitive cover $C_0 \rightarrow \tilde{E}$ in $\mathcal{H}_{N', g-2}^0(\tilde{E})$. Pick a pair of distinct points $p_1, p_2 \in C_0$ in a generic fiber of the map to \tilde{E} . Identify p_1, p_2 to get a cover $C_1 \rightarrow \tilde{E}$ with a node p . Note that the node p is \tilde{E} -superfluous, by construction.

Pick any preimages $\eta'_i \in C_1$ of $\kappa_i \in \tilde{E}$. By Proposition 5.19, $(C_1 \rightarrow \tilde{E}, \eta'_1, \eta'_2)$ is in the same component of \mathcal{P} as the original $(\tilde{C} \rightarrow E, \eta_1, \eta_2)$. Hence, if we identify η'_1 with η'_2 , we get a primitive cover $C_2 \rightarrow E$ in the same component as the original cover $C \rightarrow E$. However, $C_2 \rightarrow E$ has two nodes: η' and p . But as p was \tilde{E} -superfluous, it is E -superfluous as well. By Reduction 4, we are done.

If $g = 3$, we have to use a different argument, because there are no primitive genus 2 covers of \tilde{E} with an \tilde{E} -superfluous node. Nevertheless, we can take an isogeny $\tilde{E}_0 \rightarrow \tilde{E}$ of degree $N' - 1$, and attach an \tilde{E} tail. Let the resulting cover be $C_0 = \tilde{E}_0 \cup_p \tilde{E} \rightarrow \tilde{E}$, as in Figure 18.

Now take $\kappa_1, \kappa_2 \in \tilde{E} \subset C_0$. By Proposition 5.19, $(C_0, \kappa_1, \kappa_2)$ is in the same connected component as $(\tilde{C}, \eta_1, \eta_2)$. Let C_1 be the nodal curve obtained by identifying κ_1 and κ_2 in C_0 . We can degenerate our original cover $C \rightarrow E$ into $C_1 \rightarrow E$. But $p \in C_1$ is a disconnecting node! Hence, by Reduction 4, we are done. \square

FIGURE 18. The construction for the case $g = 3$.

Proof of Proposition 5.19. Recall that $\tilde{E} \rightarrow E$ is an isogeny, and $\mathcal{C} \rightarrow \tilde{\mathcal{H}}_{N',g-1}^0(\tilde{E})$ is the universal curve. The space $\mathcal{P} = \mathcal{C} \times_E \mathcal{C} - \Delta$ parametrizes covers $C \rightarrow \tilde{E}$ and pairs of distinct points $p_1, p_2 \in C$ that map to the same point $q \in E$ under the composite map $C \rightarrow \tilde{E} \rightarrow E$. Hence, there is a map $\mathcal{P} \rightarrow E \times \tilde{\mathcal{H}}_{N',g-1}^0(\tilde{E})$ which sends $(C \rightarrow \tilde{E}, p_1, p_2)$ to $(C \rightarrow \tilde{E}, q)$. This turns \mathcal{P} into a finite cover of an irreducible target. Hence \mathcal{P} is equidimensional, i.e., all components of \mathcal{P} have the same dimension.

Now fix a component $\Gamma \subset \tilde{E} \times_E \tilde{E}$, and let

$$\mathcal{P}_\Gamma = f^{-1}(\Gamma) \subset \mathcal{P}$$

where f is the map $f: \mathcal{P} \rightarrow \tilde{E} \times_E \tilde{E}$. We want to show that \mathcal{P}_Γ is irreducible. There is a forgetful map

$$\pi: \mathcal{P}_\Gamma \rightarrow \tilde{\mathcal{H}}_{N',g-1}^0(\tilde{E})$$

sending $(C \rightarrow \tilde{E}, p_1, p_2)$ to $(C \rightarrow \tilde{E})$. The fiber of π over $C \rightarrow \tilde{E}$ is

$$\pi^{-1}((\varphi: C \rightarrow \tilde{E})) = \{p_1, p_2 \in C, (\varphi(p_1), \varphi(p_2)) \in \Gamma \text{ and } p_1 \neq p_2\}$$

which is the same as saying that the fiber of π is the preimage of Γ in the map $C \times_E C - \Delta \rightarrow \tilde{E} \times_E \tilde{E}$. Hence, the following diagram is Cartesian.

$$\begin{array}{ccc} \pi^{-1}((\varphi: C \rightarrow \tilde{E})) & \longrightarrow & \Gamma \\ \downarrow & & \downarrow \\ C \times_E C - \Delta & \longrightarrow & \tilde{E} \times_E \tilde{E} \end{array}$$

Now observe that

- the target of π is irreducible, by Induction Hypothesis 1,
- all fibers of π have dimension one,

- the general fiber of π is irreducible: more precisely, if $C \rightarrow \tilde{E}$ is simply branched, then Lemma 4.7 says that the fiber of π is irreducible.

Moreover, we already know that \mathcal{P}_Γ is equidimensional. Hence, it follows that \mathcal{P}_Γ is irreducible. \square

5.3. The degeneration argument. Given Final Reduction 1, to complete the proof of Theorem 1.4 it is enough to prove the following.

Proposition 5.20. *Every component of $\tilde{\mathcal{H}}_{N,g}^0(E)$ meets the boundary Δ .*

We will use our study of hyperplane sections of $\tilde{V}_{df+Ne,g}(E)$ to produce such a cover in Δ .

Proof. Choose d large enough so that the Severi variety $\tilde{V}_{df+Ne,g}(E)$ dominates the Hurwitz space $\tilde{\mathcal{H}}_{N,g}(E)$. For example, by Riemann–Roch $d = g + 3$ is enough. For any component of $\tilde{\mathcal{H}}_{N,g}^0(E)$, there is a component of $\tilde{V}_{df+Ne,g}(E)$, say V_0 , dominating it. Now intersect V_0 with multiple hyperplanes of the form H_p , as in Section 3. Eventually the curve $E_0 \subset E \times \mathbb{P}^1$ is split off generically. We want to see what the image of this generic point is in the Hurwitz space $\tilde{\mathcal{H}}_{N,g}^0(E) \subset \overline{\mathcal{M}}_g(E, N)$. The map is initially not defined there, but Theorem 1.20 tells us how to extend it. The resulting cover has singular source and there are no contracted components. Hence it is a point of $\Delta \subset \tilde{\mathcal{H}}_{N,g}^0(E)$. \square

This finishes the proof of Theorem 1.4.

Let us mention two alternative ways to establish Proposition 5.20 in a more topological vein. To produce a node in the source, we have to bring a pair branch points together, in a way that, just before they meet, their monodromy transpositions exchange the same pair of sheets. To ensure this, we have to choose the path of approach carefully.

If one keeps track of how the permutation data changes as the branch points move around, one can re-encode the question in a combinatorial setting—we want to show that from any initial data, there is a sequence of moves that will result in a pair of adjacent branch points with the same monodromy transposition. However, it is still a tough problem. For example, in [GHS02, Kan05b], they can exhibit such a sequence of moves only when there are at least $2d - 2$ branch points.

In [BE84], Berstein and Edmonds ask the same combinatorial problem, but allowing not only the moves on the branch points, but also Dehn twists on the target. In Proposition 5.2 of [BE84], they show that under this setup, whenever the target has genus one, we can apply a sequence of moves to produce a pair of adjacent branch points with the same monodromy. Tracing back through their argument, we see that no Dehn twists are used, so their result holds in our setup as well, effectively proving Proposition 5.20.

As the genus of the target increases, the combinatorial problem gets even harder. Gabai and Kazez get around this challenge with the following beautiful construction.

Lemma 5.21 (Gabai–Kazez). *For an arbitrary smooth curve D , every component of $\tilde{\mathcal{H}}_{d,g}(D)$ meets the boundary Δ .*

This is similar to Lemma 3.5 in [GK87].

Proof. Given a cover $C \rightarrow D$, we want to choose a path $\gamma : [0, 1] \rightarrow D$ connecting a pair of branch points, in a way that the pair of sheets exchanged by the ramification point over $\gamma(0)$ get transported to the pair of sheets exchanged over $\gamma(1)$.

Choose a generic map $h : C \rightarrow [0, 1]$. If we perturb h , then $C \rightarrow D \times [0, 1]$ is a submersion, and the image has at most double curves and some triple points. The double curves can be closed curves, or have endpoints at the ramification points of $C \rightarrow D$. Conversely, at every ramification point there will be exactly one double curve ending there.

Take a double curve connecting a pair of ramification points, and let $\gamma : [0, 1] \rightarrow D$ be its image in D . The sheets exchanged over the branch points $\gamma(0)$ and $\gamma(1)$ are exactly the two sheets that meet to form the double curve. So the transpositions agree, and we are done! \square

5.4. Proof of Lemma 5.14. Let us recall the statement we want to prove. We are given a reduced isogeny $\pi : \tilde{E} \rightarrow E$ and an integer $D \geq 2$. We want to construct a reduced degree D isogeny $\hat{\pi} : \hat{E} \rightarrow E$ and a pair of points $p_1, p_2 \in \hat{E}$ such that

- (1) There is no common non-trivial subcover to π and $\hat{\pi}$.
- (2) For any subcover $\pi' : \hat{E} \rightarrow E'$, the points $\pi'(p_1), \pi'(p_2) \in E'$ are the same if and only if $\pi' = \hat{\pi}$ (that is, the subcover is trivial).

We are going to translate this into the language of lattices. Let $\Lambda \subset \mathbb{C}$ be the rank 2 lattice such that $E = \mathbb{C}/\Lambda$. Our isogeny $\tilde{E} \rightarrow E$ corresponds to a sublattice $\tilde{\Lambda} \subset \Lambda$. To construct the second cover $\hat{E} \rightarrow E$, we need to find another sublattice $\hat{\Lambda} \subset \Lambda$ such that

- (0') $\#(\Lambda/\hat{\Lambda}) = D$ (this corresponds to the degree of $\hat{E} \rightarrow E$ being D).
- (1') $\tilde{\Lambda} + \hat{\Lambda} = \Lambda$ (and this to having no non-trivial common subcover)

Given two points p_1, p_2 in a fiber of $\hat{E} \rightarrow E$, we can pick a path between them, which produces an element $v \in \Lambda$. Conversely, given $v \in \Lambda$ and any point $p_1 \in \hat{E}$, we can lift v and get a second point $p_2 \in \hat{E}$. Condition (2) above translates to

- (2') $\hat{\Lambda} + \mathbb{Z}v = \Lambda$

Let m be the largest integer such that $\tilde{\Lambda} \subset m\Lambda$. Note that saying that $\tilde{E} \rightarrow E$ is reduced means that $m = 1$.

Proposition 5.22. *There exist a choice of $\hat{\Lambda}$ and v satisfying conditions (0'), (1') and (2') if and only if $\gcd(D, m) = 1$.*

The case $m = 1$ implies Lemma 5.14.

Proof. We identify Λ with $\mathbb{Z} \oplus \mathbb{Z}$ by choosing a basis.

Lemma 5.23. *There is $(am, bm) \in \tilde{\Lambda}$ such that $\gcd(a, b) = 1$.*

Proof. Applying the lemma to $\tilde{\Lambda}/m$, we reduce to the case $m = 1$. Pick the smallest positive d such that $(0, d) \in \tilde{\Lambda}$. Choose another vector $(a, b) \in \Lambda$ such that (a, b) and $(0, d)$ generate Λ . In particular $\gcd(a, b, d) = 1$.

It may be that $\gcd(a, b) \neq 1$. But we may temper with it by adding arbitrary multiples of $(0, d)$. So it is enough to find an n for which $\gcd(a, b + nd) = 1$. For

each prime p dividing a , we have that p does not divide d or b , and therefore there is a residue class n_p modulo p such that

$$n \equiv n_p \pmod{p} \implies p \nmid b + nd$$

Using the Chinese remainder theorem, we can find an integer n such that no prime dividing a divides $b + nd$. Hence, $\gcd(a, b + nd) = 1$, as we wanted to show. \square

Lemma 5.24. *Take (a', b') such that $ab' - a'b = 1$. Then the lattice*

$$\hat{\Lambda} = \mathbb{Z}(a', b') \oplus \mathbb{Z}(Da, Db)$$

satisfies $\#(\mathbb{Z} \oplus \mathbb{Z}/\hat{\Lambda}) = D$, and

$$\hat{\Lambda} + \tilde{\Lambda} = \mathbb{Z} \oplus \mathbb{Z} (= \Lambda)$$

as long as $\gcd(m, D) = 1$.

Moreover, if we take $v = (a, b)$, then

$$\hat{\Lambda} + \mathbb{Z}v = \Lambda$$

Proof. That $\#(\mathbb{Z} \oplus \mathbb{Z}/\hat{\Lambda}) = D$ follows directly from $ab' - a'b = 1$.

To compute the index of $\Lambda + \Lambda' \subset \mathbb{Z} \oplus \mathbb{Z}$, first notice that $(a, b) \in \Lambda + \Lambda'$, since both (Da, Db) and (ma, mb) are in the lattice. But (a', b') is also in $\Lambda + \Lambda'$, and (a, b) and (a', b') span $\mathbb{Z} \oplus \mathbb{Z}$. \square

Lemma 5.24 shows that if $\gcd(m, D) = 1$, then we can choose $\hat{\Lambda}, v$ satisfying the conditions of Proposition 5.22.

Conversely, suppose that $\gcd(D, m) \neq 1$. Let p be a prime dividing $\gcd(D, m)$. Look at the reduction of $(\hat{\Lambda} + \tilde{\Lambda})$ modulo p . As $\tilde{\Lambda} \subset m(\mathbb{Z} \oplus \mathbb{Z})$, we have $\tilde{\Lambda} \equiv 0 \pmod{p}$. Moreover, if (a, b) and (a', b') generate $\hat{\Lambda}$, then

$$\det \begin{pmatrix} a & b \\ a' & b' \end{pmatrix} = ab' - ba' = D \equiv 0 \pmod{p}$$

Hence (a, b) and (a', b') are dependent modulo p . Therefore, $\hat{\Lambda} + \tilde{\Lambda}$ has rank at most one in the reduction modulo p . Therefore, it can't span the full lattice! This finishes the proof of Proposition 5.22, which establishes Lemma 5.14. \square

5.5. Varying the moduli of E . We will now study the version of Hurwitz space where the target curve is allowed to vary in moduli as well (see Definition 1.6). Recall that $\mathcal{H}_{d,g,1}$ is fibered over $\mathcal{M}_{1,1}$, and the fiber over (E, p) is exactly $\mathcal{H}_{d,g}(E)$. We will use our knowledge of the components of the fiber (Corollary 5.3) to describe the components of the total space $\mathcal{H}_{d,g,1}$. More precisely, the construction in Corollary 4.5 works in pointed families, and we get a map

$$\mathcal{H}_{d,g,1} \rightarrow \bigsqcup_{d \neq \tilde{d} | d} \mathcal{H}_{\tilde{d},1,1}$$

sending a cover $(C, p) \rightarrow (E, q)$ to the maximal isogeny $(\tilde{E}, \tilde{q}) \rightarrow (E, q)$ it factors through. Corollary 5.3 says that the fibers of this map are irreducible. Hence it is enough to describe the components on $\mathcal{H}_{\tilde{d},1,1}$. That is, we want to know what are the orbits of isogenies under the monodromy action induced by varying the j -invariant of the target.

Most of the isogenies will not be fixed by this monodromy action, but there is a special family which will be: the multiplication by m map, for integers m . Hence, besides the degree \tilde{d} , there is another invariant: the maximal m such that $\tilde{E} \rightarrow E$ factors through the multiplication by m map $\times m: E \rightarrow E$. Note that $m^2 \mid \tilde{d} \mid d$.

To prove Corollary 1.10, we want to show that these are the only invariants. Here is a concrete way to prove that. The isogeny $\tilde{E} \rightarrow E$ corresponds to an inclusion of lattices, which can be represented by an inclusion

$$M : \mathbb{Z} \oplus \mathbb{Z} \rightarrow \mathbb{Z} \oplus \mathbb{Z}$$

We are allowed to act by $GL_2(\mathbb{Z})$ on the source (corresponding to change basis of the lattice for \tilde{E}). We are also allowed to act by $SL_2(\mathbb{Z})$ on the target—this is the monodromy action. Under these two actions, the 2×2 matrix M can be put in a unique standard diagonal form

$$M = \begin{bmatrix} d_1 & 0 \\ 0 & d_2 \end{bmatrix}$$

such that d_1 divides d_2 . Hence M is classified by this pair of integers (d_1, d_2) . Back to our original notation, $m = d_1$ and $\tilde{d} = \det M = d_1 d_2$. Hence these are the only two discrete invariants, as we wanted to show. This completes the proof of Corollary 1.10.

Remark 5.25. The problem of classifying $n \times m$ integer matrices modulo actions of $GL(\mathbb{Z})$ on both sides is equivalent to the classification of abelian groups. It is, as a matter of fact, one of the standard ways to derive it. See for example Theorem 6.4 in [Art91].

Remark 5.26. We could ask in general what are the components of $\mathcal{H}_{d,g,h}$. Using Theorem 1.3, we can reduce again to the case of unramified covers, that is, $g = d(h-1) + 1$. Unfortunately, for $h \geq 2$, this does not translate anymore into a simple question about lattices, as above.

5.6. Further remarks on the higher genus case. We may ask what this approach can say about Theorem 1.3, the connectedness of the Hurwitz scheme of primitive covers, in the case the target has genus greater than one. Quite a few challenges arise. To start off, we do not have the induction base case. Kani's theorem [Kan03] assured us that the Hurwitz space of primitive genus two covers of elliptic curves is connected, which got our induction kicked off. There is no known direct algebraic proof of the corresponding result when the genus of the target is at least two. Aside from this, most of the arguments in this section would go through, but not all—notably, Lemma 5.14 really used the fact that the target had genus one.

But perhaps the biggest challenge is in the analogue of Theorem 1.20. We do not know much about the Severi variety of curves in $C \times \mathbb{P}^1$, when C has genus at least two. The biggest obstruction is to get good dimension estimates on the Severi variety, and their generalizations imposing tangency conditions. Theorem 2.1 is not enough anymore. In general, the less ample the anticanonical bundle of the surface is, the harder it is to work out the deformation theory of Severi varieties.

6. COMPONENTS OF THE SEVERI VARIETY $V_{Ne+df,g}$

Let us come back to the study of Severi varieties on $E \times \mathbb{P}^1$. We want to determine what its irreducible components are. A candidate is the following.

Definition 6.1. Let $V_{Ne+df,g}^{\text{prim}} \subset |\mathcal{L}|$ be the subvariety parametrizing integral curves C of geometric genus g , such that the composition of the normalization map $\tilde{C} \rightarrow C$ with the projection map $C \rightarrow E$ is a primitive cover (that is, $\tilde{C} \rightarrow C \rightarrow E$ does not factor through non-trivial isogenies).

We conjecture the following.

Conjecture 1. *The variety $V_{Ne+df,g}^{\text{prim}}$ is irreducible, for $d \geq 3$.*

Remark 6.2. For $d = 2$ it is not clear what to expect. In view of the construction in Section 2.1.1, we see that all the nodes have to lie on four special fibers (the preimages of the four ramification points of the double cover $E \rightarrow \mathbb{P}^1$). Hence, to each component of $V_{Ne+2f,g}^{\text{prim}}$, we can associate a partition of δ in four parts (where $\delta = 2N - 1 - g$ is the number of nodes). If more than one partition occurs, we would get more than one component.

On the other hand, for $g = 2$ only one partition occurs for primitive covers (see [Kuh88]). It is not clear what the behavior will be for $g > 2$.

Using our knowledge of Hurwitz spaces, we can establish irreducibility in at least a restricted range.

Theorem 6.3. *When $d \geq 2g - 1$, the Severi variety $V_{Ne+df,g}^{\text{prim}}$ is irreducible.*

Proof. Consider the map $V_{Ne+df,g}^{\text{prim}} \rightarrow \mathcal{H}_{N,g}^0(E)$. The target is irreducible by Theorem 1.4. To establish the irreducibility of the source, it is enough to show that the map is dominant with irreducible fibers of the same dimension.

Let $C \rightarrow E$ be a point in $\mathcal{H}_{N,g}^0(E)$. The fiber over it corresponds to choosing a line bundle $L \in \text{Pic}^d(C)$, and a framing for the map $C \rightarrow \mathbb{P}^1$. The line bundle L is in the non-special range, as $d \geq 2g - 1$, and hence it has $d - g + 1$ global sections. In total, we get irreducible fibers of dimension

$$1 + 2 \times (d - g + 1) - 1$$

which is independent of the map $C \rightarrow E$. □

APPENDIX A. DEFORMATION THEORY

The goal of this section is to prove Theorem 2.1. We will only adapt the arguments already existing in the literature for our setup. In specific, we will follow [CH98] closely. Let us remind ourselves of the setup. We have

- a smooth projective surface Σ ,
- a fixed curve $D \subset \Sigma$,
- a fixed finite set $\Omega \subset D$,
- a fixed homology class $\tau \in H_2(\Sigma, \mathbb{C})$,
- and integers g and b .

We are given a substack $W \subset \mathcal{M}_g(\Sigma, \tau)$, such that the general point corresponds to maps $f: C \rightarrow \Sigma$ such that

- the source C is a smooth connected curve of genus g ,
- the map f is birational onto its image,
- the preimage $Q = f^{-1}(D - \Omega)$ is supported on at most b points.

We want to bound the dimension of W in terms of

$$\gamma = -(K_\Sigma + D) \cdot \tau + b$$

More precisely, we want to show that if $\gamma \geq 1$ then

$$\dim W \leq g - 1 + \gamma$$

and to control the singularities of $f(C) \subset \Sigma$, according to how large γ is.

To do so, we need some notation to keep track of the points in $f^{-1}(D)$. Let $q_1, \dots, q_{\bar{b}}$ be the support of $Q = f^{-1}(D - \Omega)$. By our assumption, we have $\bar{b} \leq b$. Let $p_1, \dots, p_a \in C$ be the remaining points in the support of $f^{-1}(D)$. Note that $f(p_i)$ is in the finite set Ω , and hence is not allowed to vary in W . Let α_i, β_j be such that

$$f^{-1}(D) = \sum \alpha_i p_i + \sum \beta_j q_j$$

After a finite base change, we get $\widetilde{W} \subset \mathcal{M}_{g, a+\bar{b}}(\Sigma, \tau)$ in which the points p_i, q_j become sections instead of just multisections.

We want to determine the tangent space to \widetilde{W} at f . We start by determining the tangent space to the whole Kontsevich space. Deformation theory tell us that a first order deformation of the map $f: C \rightarrow \Sigma$, without taking into account the tangency conditions, produces a section $\sigma \in H^0(C, \mathcal{N})$ where \mathcal{N} is the normal sheaf to the map f , defined by the sequence:

$$0 \rightarrow T_C \rightarrow f^* T_\Sigma \rightarrow \mathcal{N} \rightarrow 0$$

Although we expect \mathcal{N} to be torsion free, we cannot assume that from the outset. Anyways, it will be useful to consider the image $\bar{\sigma} \in H^0(\mathcal{N}/\mathcal{N}_{\text{tor}})$ of σ under $H^0(\mathcal{N}) \rightarrow H^0(\mathcal{N}/\mathcal{N}_{\text{tor}})$.

To take the tangency conditions into account, we invoke Lemma 2.6 from [CH98].

Lemma A.1 (Caporaso–Harris). *The section $\bar{\sigma}$ vanishes to order α_i at p_i , and at least $\beta_j - \ell_j$ at q_j , where $\ell_j - 1$ is the order of vanishing of df at q_j . Moreover, the exact order of vanishing of $\bar{\sigma}$ is either $\beta_j - \ell_j$ or at least β_j .*

We expect all ℓ_j to be 1, and the tangent space to \widetilde{W} to be

$$H^0\left(C, \mathcal{N}\left(-\sum \alpha_i p_i - \sum (\beta_j - 1) q_j\right)\right)$$

This sheaf has degree $2g - 2 + \gamma$, and Riemann–Roch gives $g - 1 + \gamma$ for the expected dimension of W .

This estimate can be wrong for two reasons:

- (1) The higher cohomology of $\mathcal{N}(-\sum \alpha_i p_i - \sum (\beta_j - 1) q_j)$ might not vanish.
- (2) The sheaf $\mathcal{N}(-\sum \alpha_i p_i - \sum (\beta_j - 1) q_j)$ could have torsion.

The first issue will not arise if \mathcal{N} is torsion free, and its degree is in the non-special range, which corresponds to the condition $\gamma \geq 1$.

The second issue arises when the map f has ramification (for example, if the image curve C had a cusp). Nevertheless, this problem was dealt with by Arbarello–Cornalba, and then Caporaso–Harris. They showed that for $f: C \rightarrow \Sigma$ general, the dimension of $H^0(\mathcal{N}/\mathcal{N}_{\text{tor}})$ is an upper bound on the dimension of the Kontsevich space. This assumes that the evaluation map is birational (which we are assuming already, anyways), and that the characteristic is zero.

Their setup is the following. Let $\mathcal{C} \rightarrow B$ be a flat, smooth and proper family of curves over a smooth connected base B , and $f: \mathcal{C} \rightarrow \Sigma \times B$ be a map over B . To each fiber $f_b: C_b \rightarrow \Sigma$, we associate the normal sheaf \mathcal{N}_b for f_b . Then the *Horiwaka map* sends a tangent vector at $b \in B$ to the corresponding first order deformation of the map f_b :

$$k_b: T_b B \rightarrow H^0(C_b, \mathcal{N}_b)$$

If the family $\mathcal{C} \rightarrow \Sigma \times B$ is nowhere isotrivial, this map is injective, which gives an upper bound for the dimension of B .

Lemma A.2 (Arbarello–Cornalba, Caporaso–Harris). *If $\mathcal{C} \rightarrow S \times B$ is birational onto its image, then for $b \in B$ general*

$$\text{im } k_b \cap H^0(C_b, (\mathcal{N}_b)_{\text{tor}}) = 0$$

We use this lemma to deal with the issues raised above, and prove Theorem 2.1.

Proof of Theorem 2.1. Consider the following divisors on C :

$$\begin{aligned} D &= \sum \alpha_i p_i + \sum \beta_j q_j \\ D_0 &= \sum \ell_j q_j \\ D_1 &= D - D_0 \end{aligned}$$

Note that, by an abuse of notation, $D = f^*(\mathcal{O}_\Sigma(D))$.

By Lemma A.1, the image of $T_f W \subset H^0(C, \mathcal{N})$ on $H^0(C, \mathcal{N}/\mathcal{N}_{\text{tor}})$ is contained in $H^0(C, \mathcal{N}/\mathcal{N}_{\text{tor}}(-D_1))$. And by Lemma A.2, the dimension of the latter is an upper bound on the dimension of W . That is,

$$\dim W \leq h^0(C, \mathcal{N}/\mathcal{N}_{\text{tor}}(-D_1))$$

We can describe very explicitly what the line bundle $\mathcal{N}/\mathcal{N}_{\text{tor}}(-D_1)$ is. The sheaf \mathcal{N}_{tor} is supported on the ramification points of $f: C \rightarrow \Sigma$. Let Γ be the skyscraper sheaf of ramification away from the q_j . Then

$$\mathcal{N}_{\text{tor}} = \Gamma \oplus \mathbf{k}_{q_j}^{\ell_i-1}$$

And, as we know that

$$c_1(\mathcal{N}) = K_C - f^*K_\Sigma$$

we get

$$\begin{aligned} \mathcal{N}/\mathcal{N}_{\text{tor}} &= K_C \otimes f^*K_\Sigma^{-1}(-c_1(\Gamma) - \sum (\ell_i - 1)q_j) \\ &= K_C(-f^*K_\Sigma - D_0 + \sum q_j - c_1(\Gamma)) \end{aligned}$$

Hence,

$$\mathcal{N}/\mathcal{N}_{\text{tor}}(-D_1) = K_C(-f^*K_\Sigma - D + \sum q_j - c_1(\Gamma))$$

which is a sub sheaf of

$$\mathcal{L} = K_C(-f^*K_\Sigma - D + \sum q_j)$$

whose degree is

$$\deg \mathcal{L} = 2g - 2 - (K_\Sigma + D)\tau + \bar{b}$$

Adding an effective divisor B of degree $b - \bar{b}$, we get that $\deg \mathcal{L}(B) = 2g - 2 + \gamma$, and $\mathcal{N}/\mathcal{N}_{\text{tor}}(-D_1)$ is a subsheaf of $\mathcal{L}(B)$.

Hence, if $\gamma \geq 1$, we are in the non-special range, and Riemann–Roch tells us

$$\begin{aligned} \dim W &\leq h^0(\mathcal{N}/\mathcal{N}_{\text{tor}}(-D_1)) \\ &\leq h^0(\mathcal{L}(B)) = g - 1 + \gamma \end{aligned}$$

Let us assume from now on that equality holds.

If $\gamma \geq 2$, then $\Gamma = \emptyset$ and $b = \bar{b}$. Indeed, if $r \in \text{Supp } \Gamma$, or in $\text{Supp } B$, then

$$\dim W \leq h^0(\mathcal{L}(B - r))$$

which is still in the non-special range, and hence gives too small a dimension.

To be unramified along the q_j we need more freedom—let us assume that $\gamma \geq 3$. Then we can find a section of \mathcal{L} vanishing to order exactly 1 on q_j . Hence, the order of vanishing of $\bar{\sigma} \in H^0(\mathcal{N}/\mathcal{N}_{\text{tor}})$ at q_j is $\beta_j - \ell_j + 1$. By the Lemma A.1, this must be at least β_j , and hence $\ell_j = 1$. Hence f is unramified and \mathcal{N} is torsion-free. Therefore, $\mathcal{N} = K_C(-f^*K_\Sigma)$.

To check smoothness of $f(C)$ along D , we only need show that pairs of points q_j don't always map to the same point. We ensure that by finding a section vanishing at q_{j_1} but not at q_{j_2} . Again, $\gamma \geq 3$ is enough.

We can rule out tacnodes too. If r_1, r_2 map down to a tacnode, we need to find a section vanishing at one but not the other, which we are guaranteed to find, again by Riemann–Roch.

Finally, let us rule out triple points. Suppose there is one, and that $r_1, r_2, r_3 \in C$ are its preimages. We would like to find a section vanishing at r_1 and r_2 , but not r_3 . To guarantee that we need $\mathcal{N}(-D + \sum q_j - r_1 - r_2 - r_3)$ to be non-special, which holds when $\gamma \geq 4$ \square

REFERENCES

- [AP05] Michela Artebani and Gian Pietro Pirola. Algebraic functions with even monodromy. *Proceedings of the American Mathematical Society*, 133(2):331–341 (electronic), 2005.
- [Art91] Michael Artin. *Algebra*. Prentice Hall, Inc., Englewood Cliffs, NJ, 1991.
- [BE84] Israel Berstein and Allan L. Edmonds. On the classification of generic branched coverings of surfaces. *Illinois Journal of Mathematics*, 28(1):64–82, March 1984.
- [CH98] Lucia Caporaso and Joe Harris. Counting plane curves of any genus. *Inventiones mathematicae*, 131(2):345–392, 1998.
- [Che10] Dawei Chen. Covers of elliptic curves and the moduli space of stable curves. *Journal für die Reine und Angewandte Mathematik. [Crelle's Journal]*, 649:167–205, 2010.
- [Cil99] Ciro Ciliberto. Nodal curves on surfaces. <http://mat.uab.es/~kock/RLN/ciro.ps>, 1999. [Lecture notes by Joachim Kock; online; accessed 25-May-2014].

- [Cle73] Alfred Clebsch. Zur Theorie der Riemann'schen Fläche. *Mathematische Annalen*, 6(2):216–230, 1873.
- [CP12] Yaim Cooper and Rahul Pandharipande. A Fock space approach to Severi degrees. *arXiv:1210.8062 [math]*, 2012.
- [DH88] Steven Diaz and Joe Harris. Ideals associated to deformations of singular plane curves. *Transactions of the American Mathematical Society*, 309(2):433–468, 1988.
- [FM01] Daniel Frohardt and Kay Magaard. Composition factors of monodromy groups. *Annals of Mathematics. Second Series*, 154(2):327–345, 2001.
- [Ful69] William Fulton. Hurwitz schemes and irreducibility of moduli of algebraic curves. *Annals of Mathematics. Second Series*, 90:542–575, 1969.
- [GHS02] Tom Graber, Joe Harris, and Jason Starr. A note on Hurwitz schemes of covers of a positive genus curve. *arXiv:math/0205056*, 2002.
- [GK87] David Gabai and William H. Kazez. The classification of maps of surfaces. *Inventiones Mathematicae*, 90(2):219–242, 1987.
- [GMS03] Robert M. Guralnick, Peter Müller, and Jan Saxl. The rational function analogue of a question of Schur and exceptionality of permutation representations. *Memoirs of the American Mathematical Society*, 162(773):viii+79, 2003.
- [Gur03] Robert Guralnick. Monodromy groups of coverings of curves. In *Galois groups and fundamental groups*, volume 41 of *Math. Sci. Res. Inst. Publ.*, pages 1–46. Cambridge Univ. Press, Cambridge, 2003.
- [Har86] Joe Harris. On the Severi problem. *Inventiones Mathematicae*, 84(3):445–461, 1986.
- [HM82] Joe Harris and David Mumford. On the Kodaira dimension of the moduli space of curves. *Inventiones Mathematicae*, 67(1):23–88, 1982. With an appendix by William Fulton.
- [HM98] Joe Harris and Ian Morrison. *Moduli of Curves*, volume 187 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998.
- [Kan03] Ernst Kani. Hurwitz spaces of genus 2 covers of an elliptic curve. *Universitat de Barcelona. Collectanea Mathematica*, 54(1):1–51, 2003.
- [Kan04] Vassil Kanev. Hurwitz spaces of triple coverings of elliptic curves and moduli spaces of abelian threefolds. *Annali di Matematica Pura ed Applicata. Series IV*, 183(3):333–374, 2004.
- [Kan05a] Vassil Kanev. Hurwitz spaces of quadruple coverings of elliptic curves and the moduli space of abelian threefolds $\mathcal{A}_3(1, 1, 4)$. *Mathematische Nachrichten*, 278(1-2):154–172, 2005. *arXiv:math/0308006*.
- [Kan05b] Vassil Kanev. Irreducibility of Hurwitz spaces. *arXiv:math/0509154*, 2005.
- [Kan13] Vassil Kanev. Unirationality of hurwitz spaces of coverings of degree ≤ 5 . *International Mathematics Research Notices. IMRN*, (13):3006–3052, 2013.
- [Kuh88] Robert M. Kuhn. Curves of genus 2 with split Jacobian. *Transactions of the American Mathematical Society*, 307(1):41–49, 1988.
- [SS13] Mendy Shoval and Eugenii Shustin. On Gromov–Witten invariants of Del Pezzo surfaces. *International Journal of Mathematics*, 24(07):1350054, 2013.
- [Ty07] Ilya Tyomkin. On Severi varieties on Hirzebruch surfaces. *International Mathematics Research Notices. IMRN*, (23):Art. ID rnm109, 31, 2007.
- [Vak00] Ravi Vakil. Counting curves on rational surfaces. *Manuscripta Mathematica*, 102(1):53–84, 2000.
- [Vis97] Angelo Vistoli. The deformation theory of local complete intersections. *arXiv:alg-geom/9703008*, 1997.

HARVARD MATHEMATICS DEPARTMENT, 1 OXFORD STREET, CAMBRIDGE, MA
 E-mail address: gbujokas@math.harvard.edu